IPv6 Technical Overview

Andrew Yang (andryang@cisco.com)

System Engineer. Sales

Cisco Taiwan

© 2011 Cisco and/or its affiliates. All rights reserved.

Agenda

- Why IPv6
- What is IPv6
- IPv6 Addressing part 1
- IPv6 Addressing part 2
- The role of Multicast
- IPv6 Interface Configurations
- Fragmentation and PMTU
- ICMPv6 and Neighbor Discovery

© 2011 Cisco and/or its affiliates. All rights reserved.

Why IPv6

© 2011 Cisco and/or its affiliates. All rights reserved.

RIR Pool Exhaustion Dates

- RIRs still have addresses available from local pools
- APNIC, RIPE, and ARIN will all be exhausted by 2012

Nortel sold 666,624 IP addresses to Microsoft on 23rd March for $7.5M ($11.25 each)

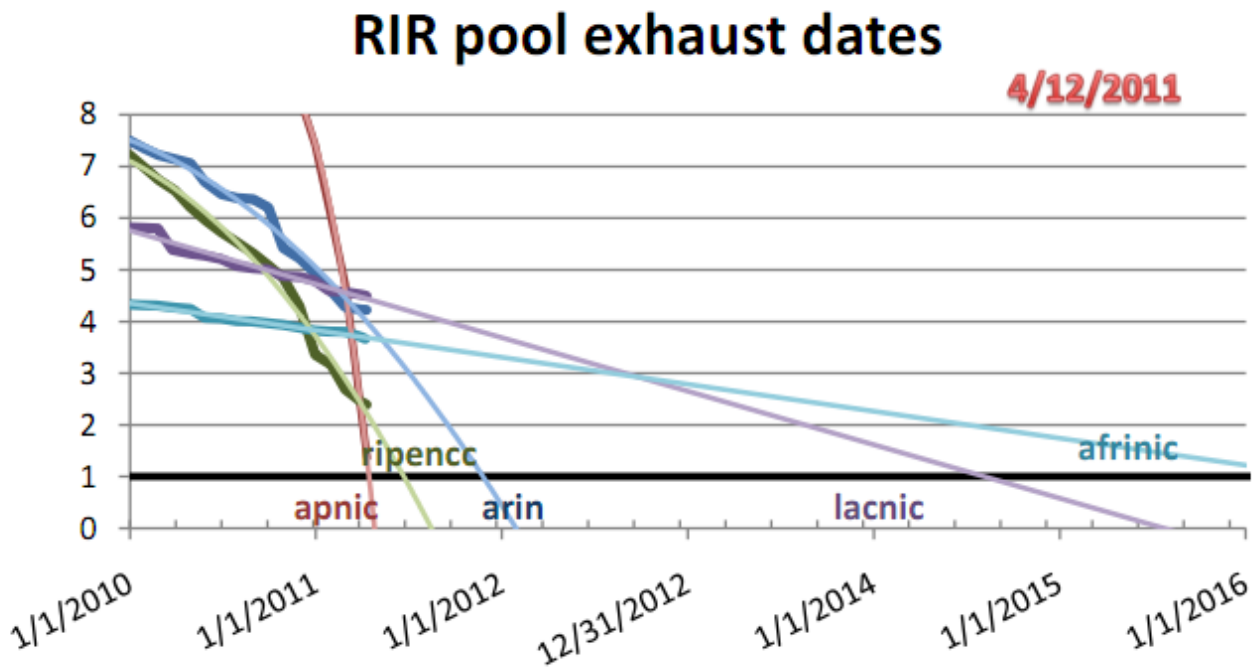

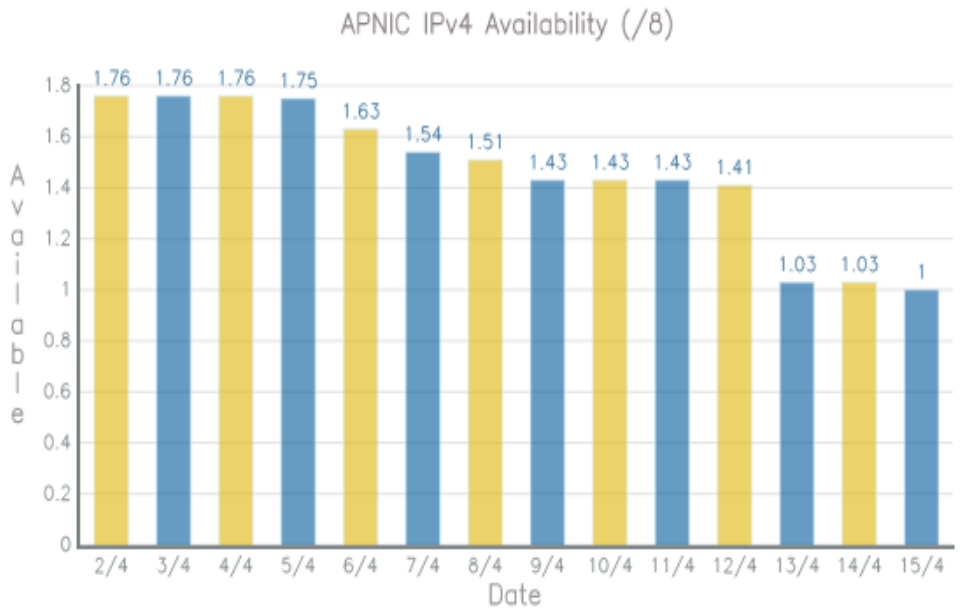

Source: Tony Hain, Cisco Systems http://www.tndh.net/~tony/ietf/IPv4-rir-pools.pdf

© 2011 Cisco and/or its affiliates. All rights reserved.

Failure to Act will Impact Business

2010 | 2012 | 2014

- 2010: Low Impact – Buying behavior shift limited to mandated and early adopter sites

Transition Planning

Early Adopters

Globalization

IPv6 Government Mandate Deadlines

IPv4/IPv6 Co-existence

2011: Internet Evolution begins – "*…IPv6 is important to all of us (…) to everyone around the world, It is crucial to our ability to tie together everyone and every device*". *John Chambers*

2012: Mandates take effect – Transition to IPv6 forces customers to acquire product or managed services to sustain business and customer reach

2014: IPv6 is mainstream – customers without transition infrastructure experience reduced service levels, diminished customer reach, increase operational complexity

IPv6 Business Impact – The Cost of Waiting Goes Up

Low Risk | Moderate Risk | High Risk

IPv6 Global Deployment to Users

10+ Years for Preparation and Trial..

Now Doubling every 9 months!

Government Mandates have played a significant role in deployment.

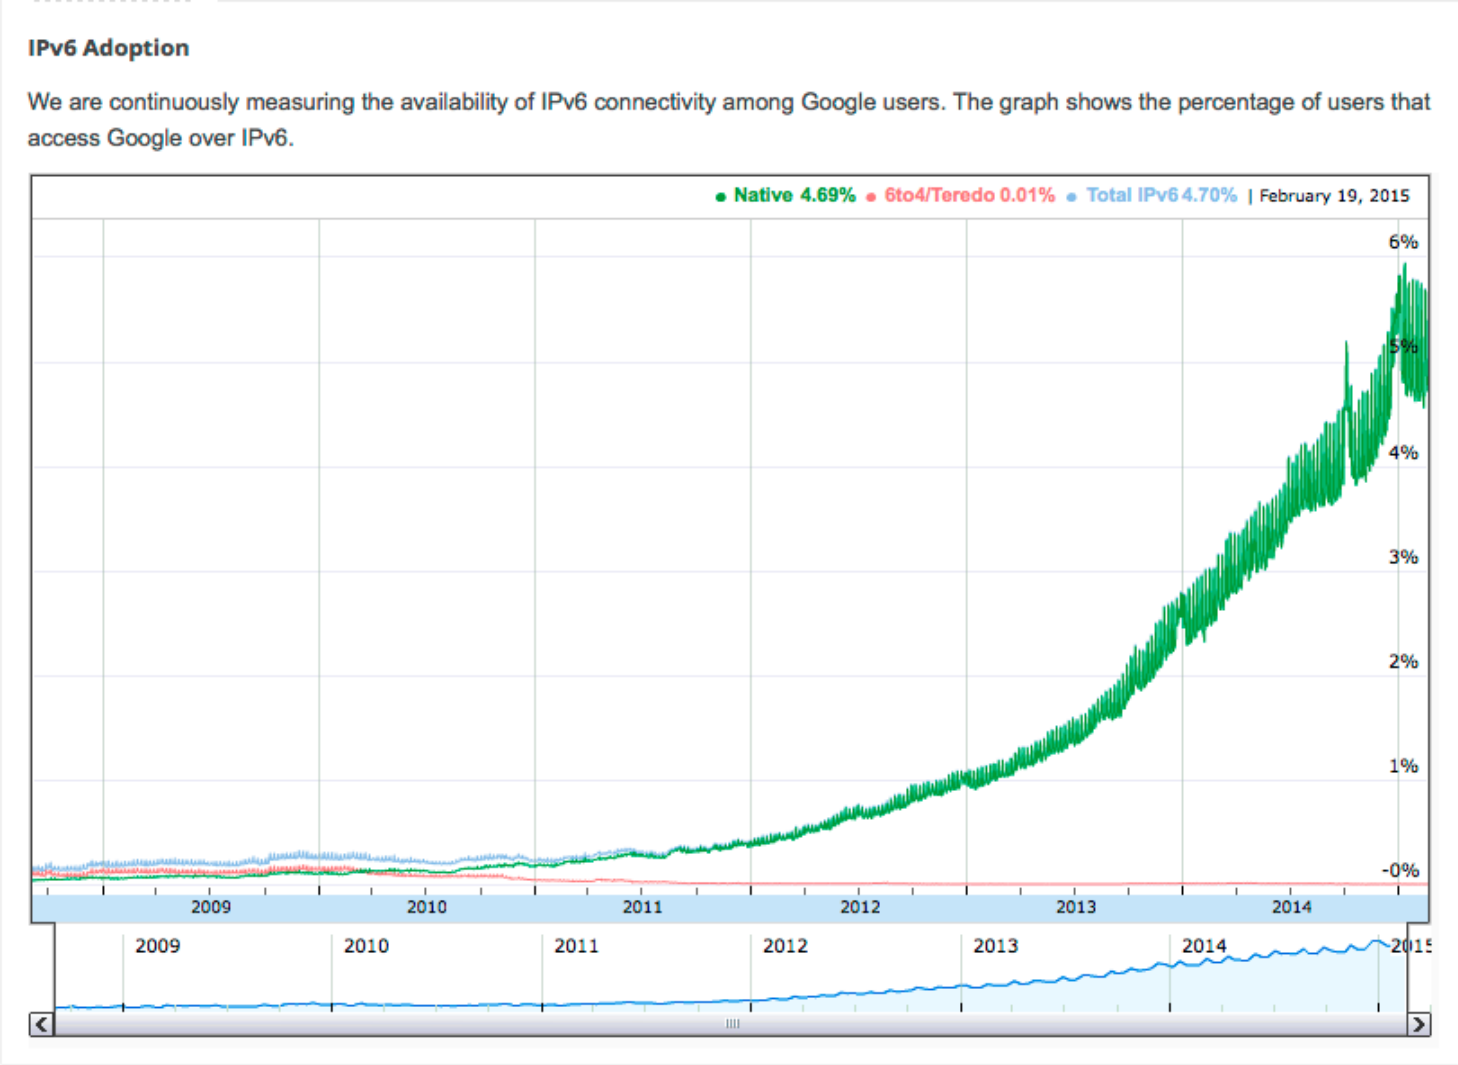

https://www.google.com/intl/en/ipv6/statistics.html

© 2011 Cisco and/or its affiliates. All rights reserved.

Connected Devices Growth

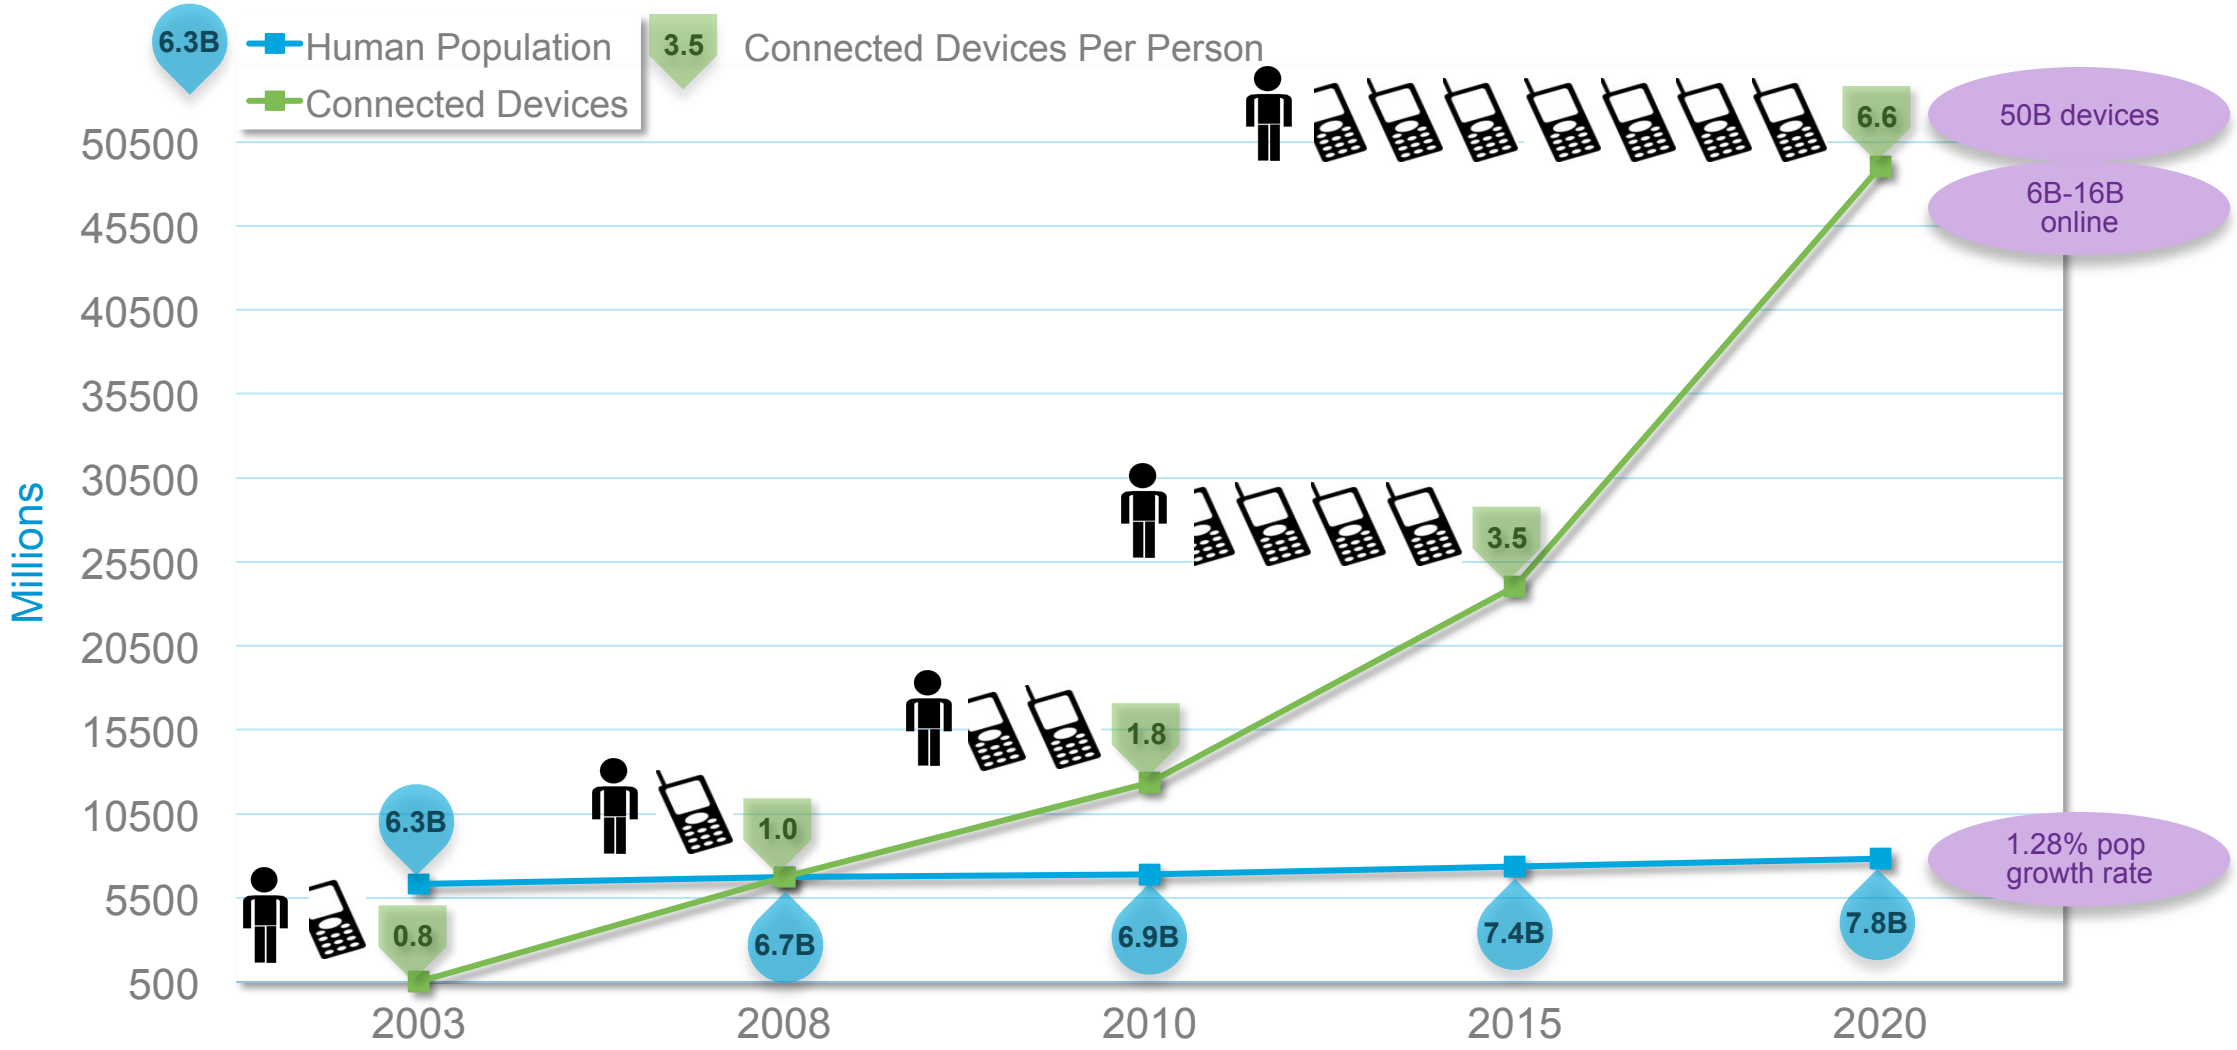

Cisco IBSG projections, UN Economic & Social Affairs http://www.un.org/esa/population/publications/longrange2/WorldPop2300final.pdf

© 2011 Cisco and/or its affiliates. All rights reserved.

IPv6 Drivers

IPv6 is an enabler
It is NOT a new service
It allows anything to connect to everything

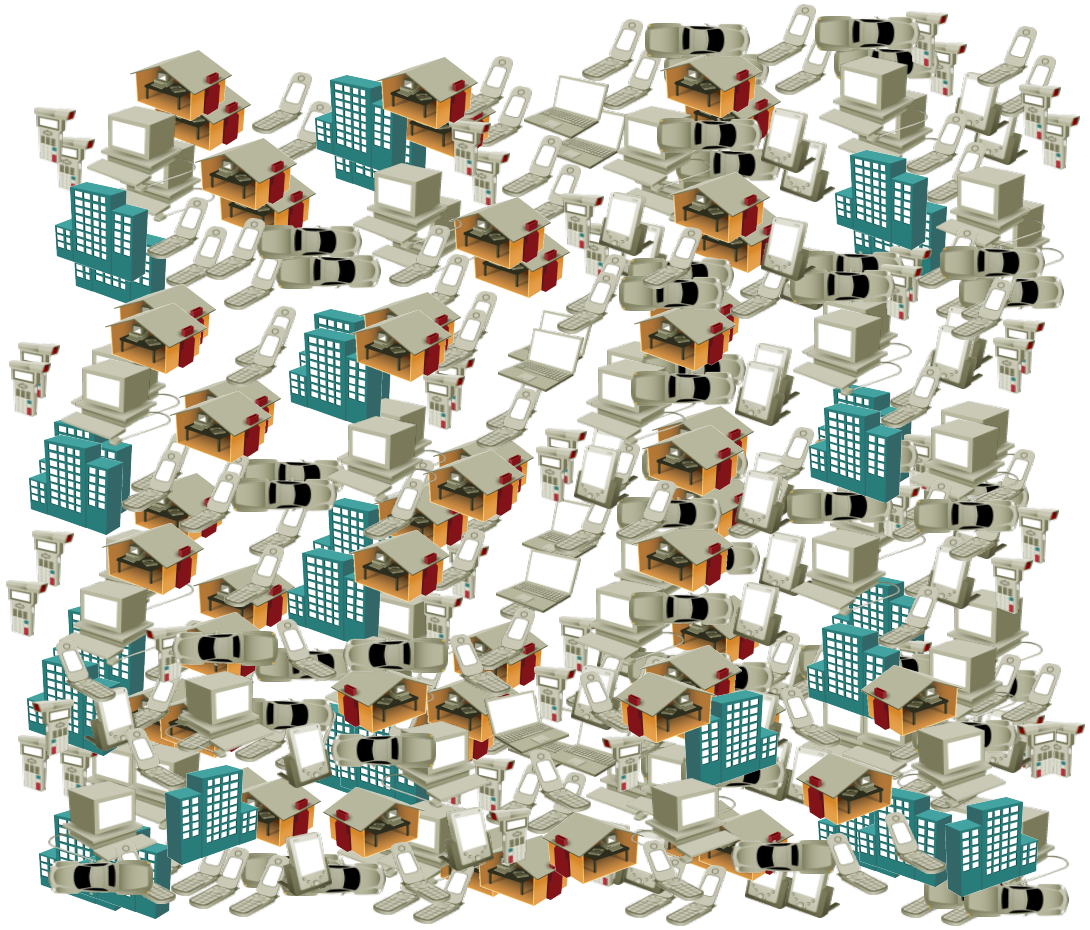

- Mandatory
  - Government Mandates
  - IPv4 address pool exhausted
  - NGN Capabilities to Defence
- Large Address Consumption Requirements
  - Cable market address scaling
  - Population densities in APAC
  - 4G deployments
- IPv6 Technologies
  - Smart Grids/Sensor Networks/6LowPAN
  - Connected Communities
- IPv4 connects computers
- IPv6 connects people and things

© 2011 Cisco and/or its affiliates. All rights reserved.

An Application of IPv6 and Smart Grids

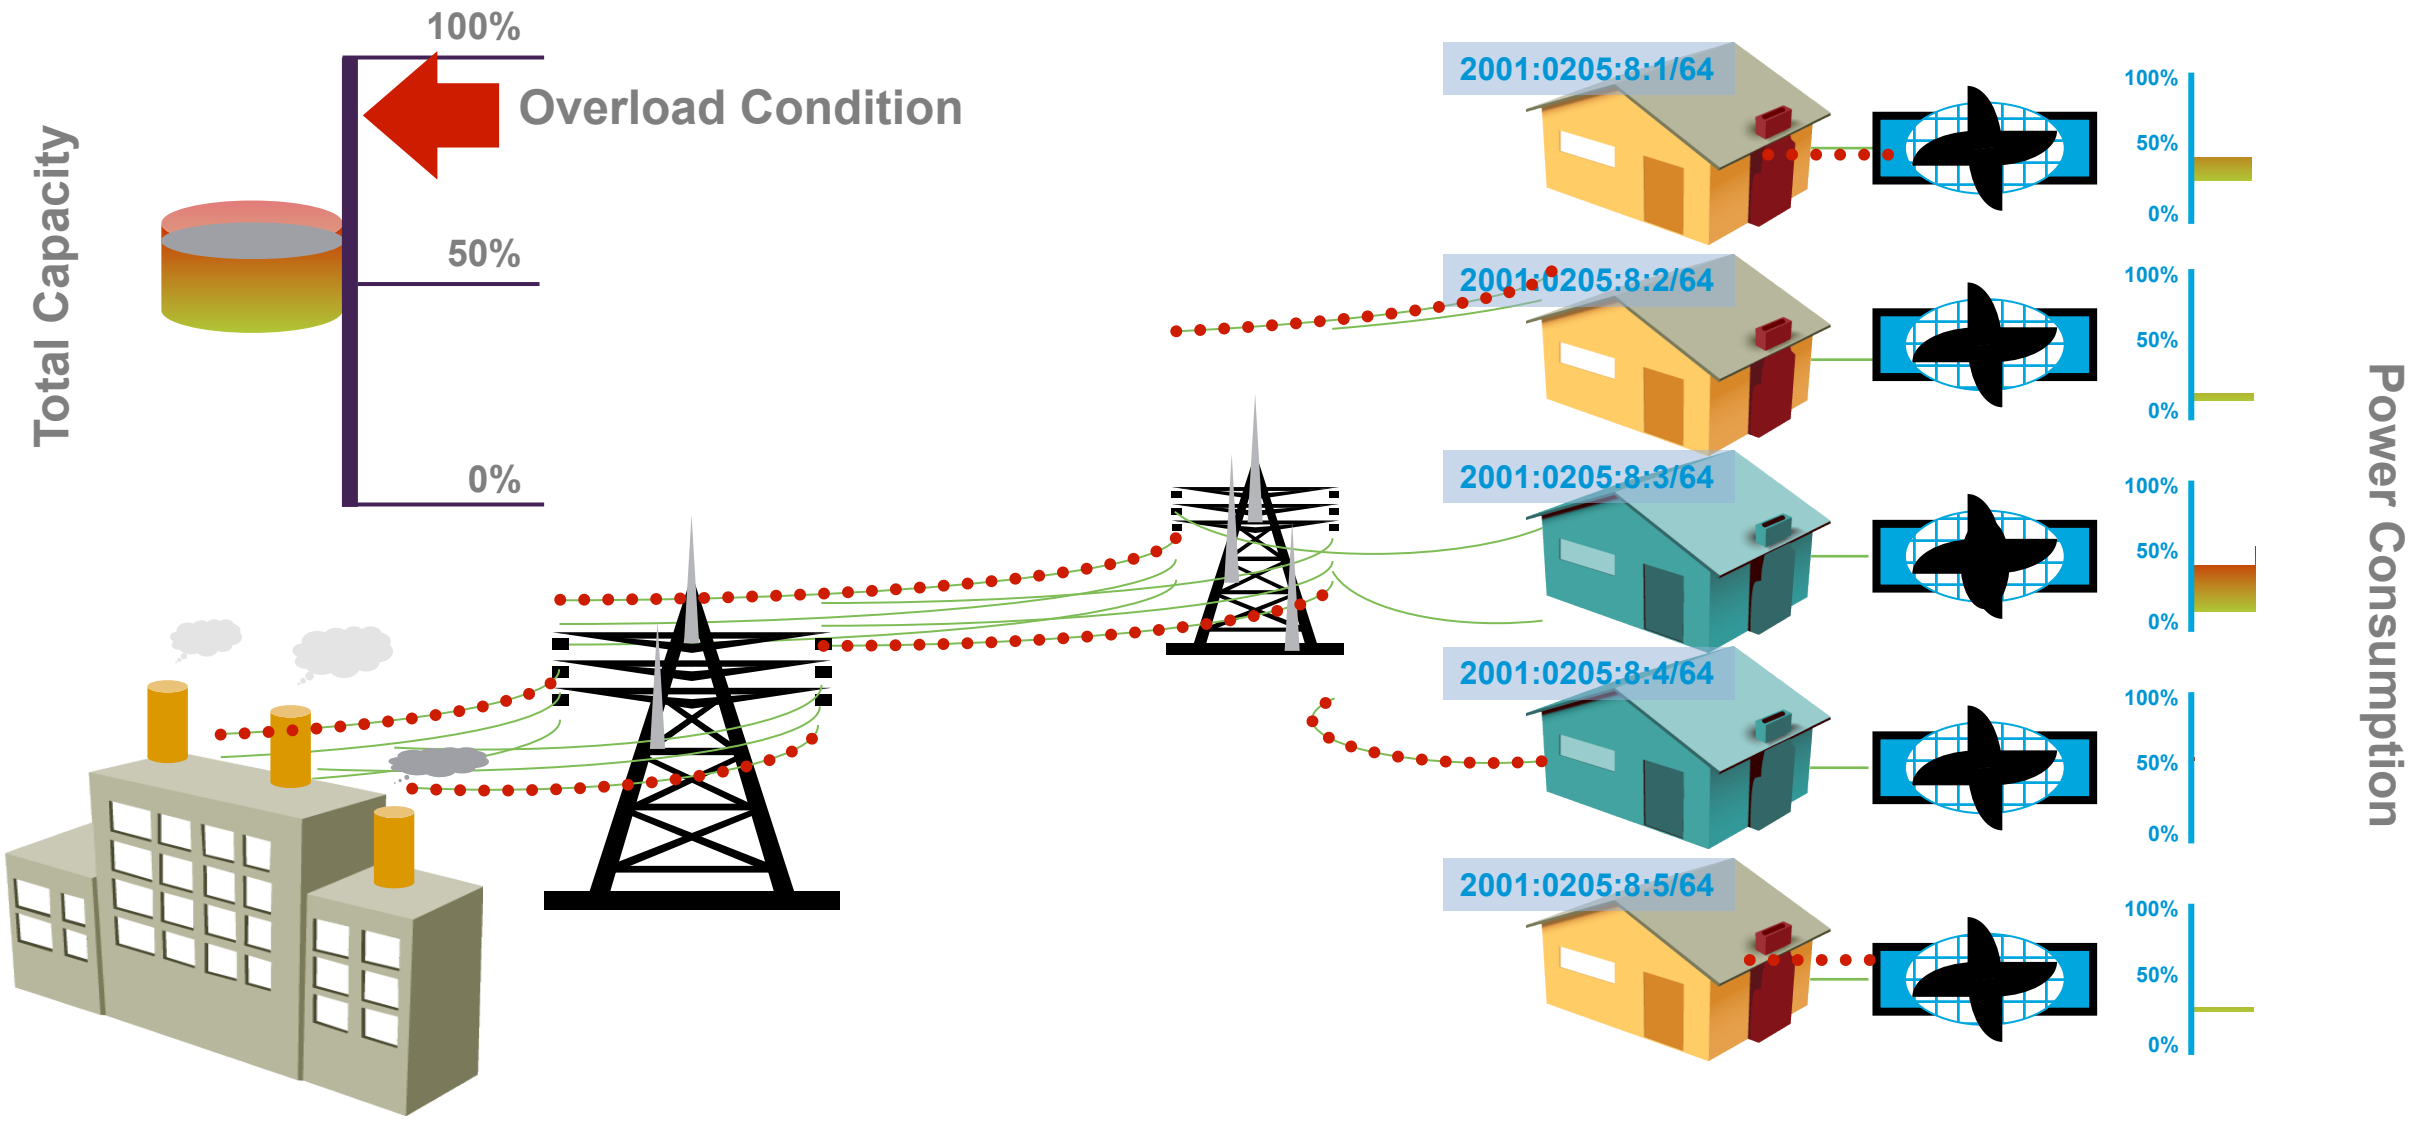

© 2011 Cisco and/or its affiliates. All rights reserved.

Home Networking & IPv6

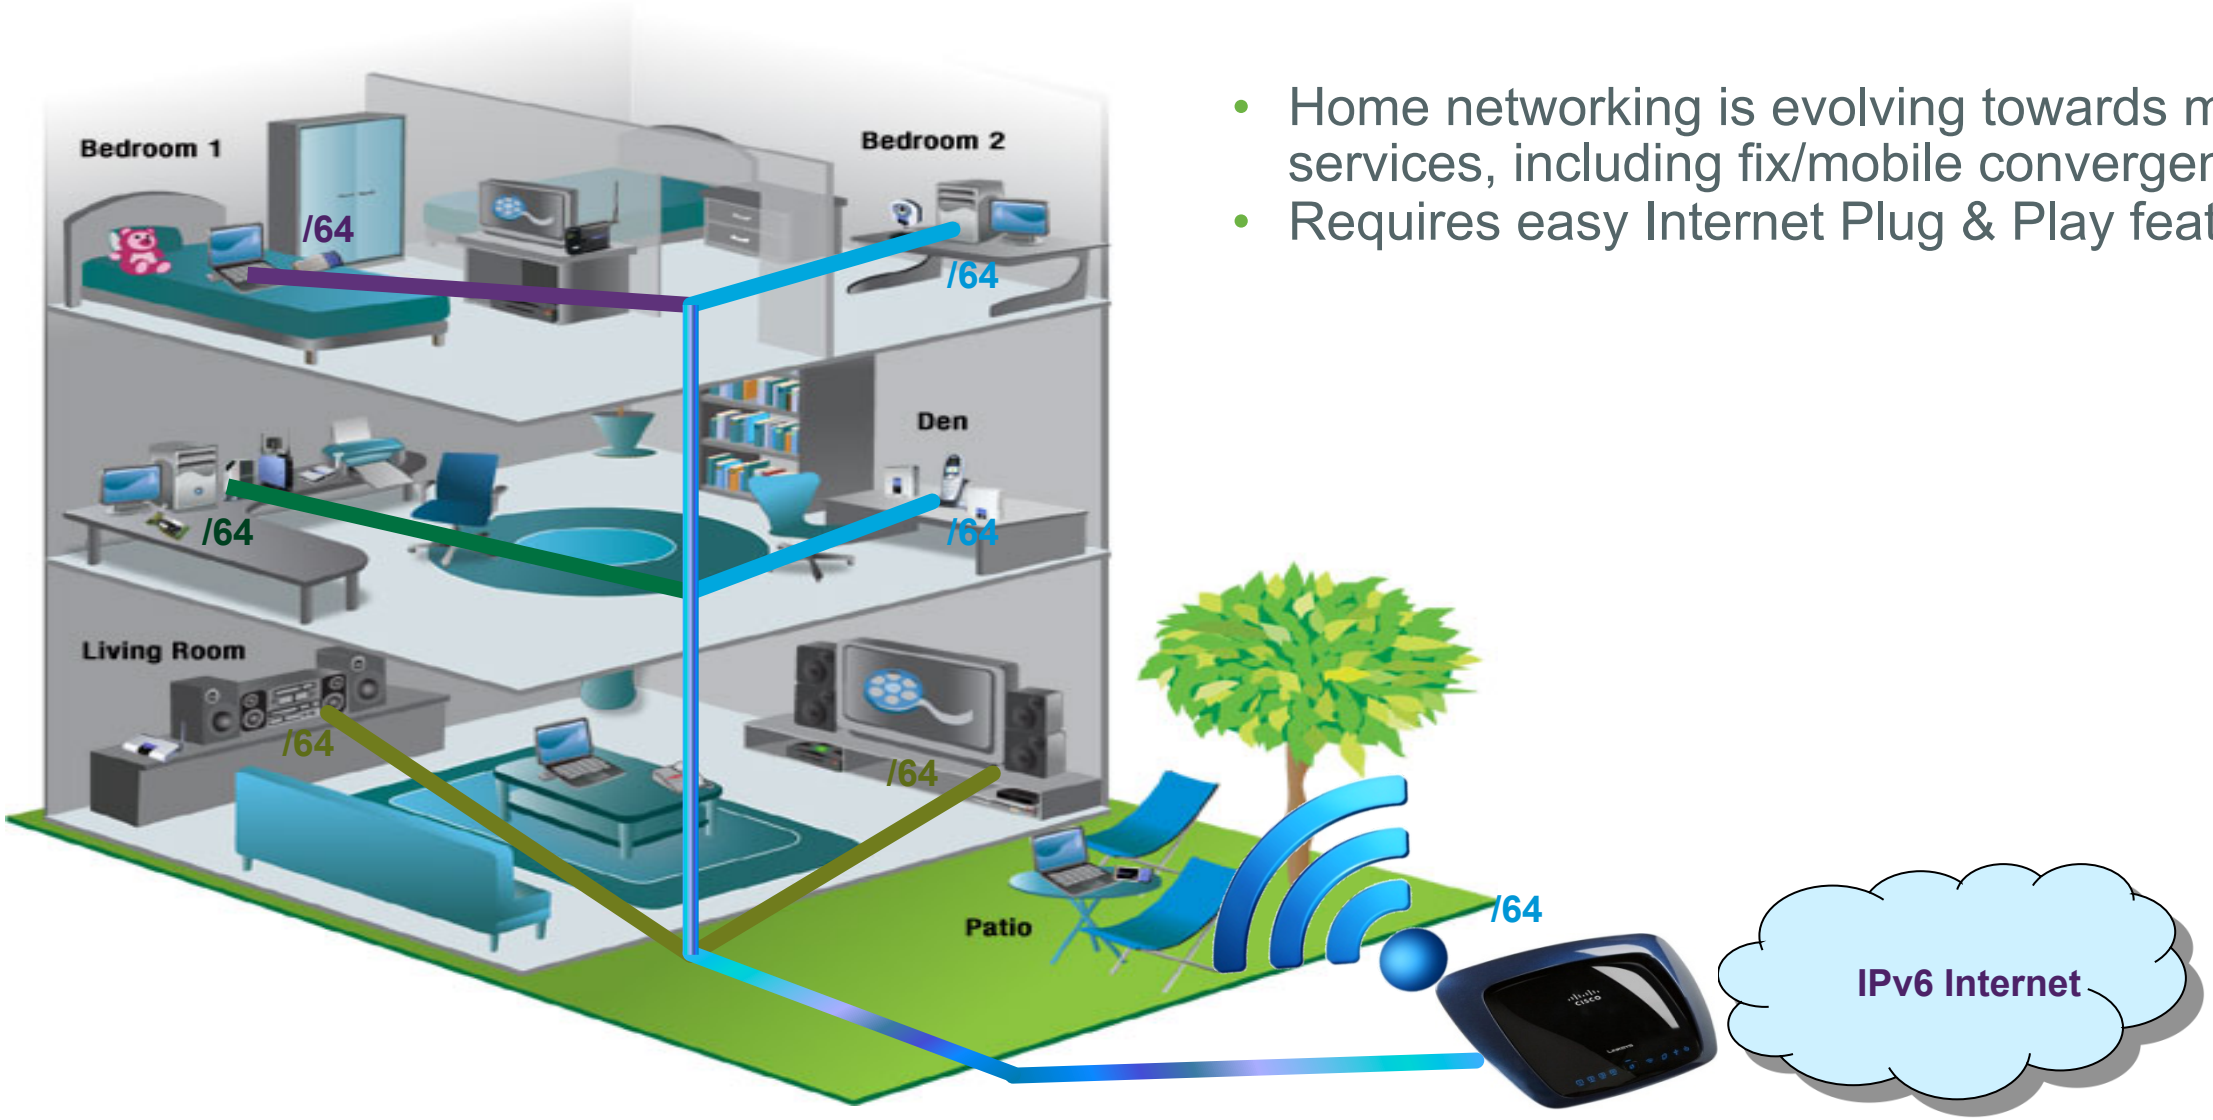

- Home networking is evolving towards more services, including fix/mobile convergence
- Requires easy Internet Plug & Play features

© 2011 Cisco and/or its affiliates. All rights reserved.

What is IPv6

© 2011 Cisco and/or its affiliates. All rights reserved.

What is IPv6

- Defined by RFC2460 (1998)
- 128bit addressing scheme to replace the 32bit IPv4 addressing scheme
  - Hexadecimal representation
  - CIDR (Classless Inter-Domain Routing) masking only
- Introduces new protocol level behaviours
  - Neighbour Discovery
  - Stateless Addressing
  - No more Broadcast, only Multicast

© 2011 Cisco and/or its affiliates. All rights reserved.

So How Big Is The IPv6 Address Space?

340,282,366,920,938,463,463,374,607,432,768,211,456

(IPv6 Address Space - 340 Trillion Trillion Trillion)

vs

4,294,967,296

(IPv4 Address Space - 4 Billion)

© 2011 Cisco and/or its affiliates. All rights reserved.

So How Big Is The IPv6 Address Space?

Antares
15th Brightest
star in the sky

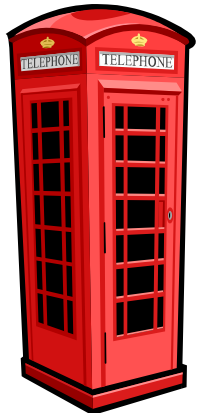

- Let's assume our Sun represents 4 Billion Addresses
- The IPv6 Address space would approach the size Antares
- In fact, a proper comparison would be to compare Antares with a Telephone Box

© 2011 Cisco and/or its a

IPv6 over Ethernet

- IPv6 uses Ethernet Protocol ID (0x86DD)

Dest MAC	Source MAC	0x86DD	IPv6 Header and Payload

- IPv4 uses Ethernet Protocol ID (0x0800)

Dest MAC	Source MAC	0x0800	IPv4 Header and Payload

© 2011 Cisco and/or its affiliates. All rights reserved.

IPv4 and IPv6 Header Comparison

IPv4 Header

Version	IHL	Type of Service	Total Length	
Identification			Flags	Fragment Offset
Time to Live		Protocol	Header Checksum	
Source Address				
Destination Address				
Options				Padding

IPv6 Header

Version	Traffic Class	Flow Label		
Payload Length			Next Header	Hop Limit
Source Address				
Destination Address				

Legend

- Field's Name Kept from IPv4 to IPv6
- Fields Not Kept in IPv6
- Name and Position Changed in IPv6
- New Field in IPv6

© 2011 Cisco and/or its affiliates. All rights reserved.

Extension Headers

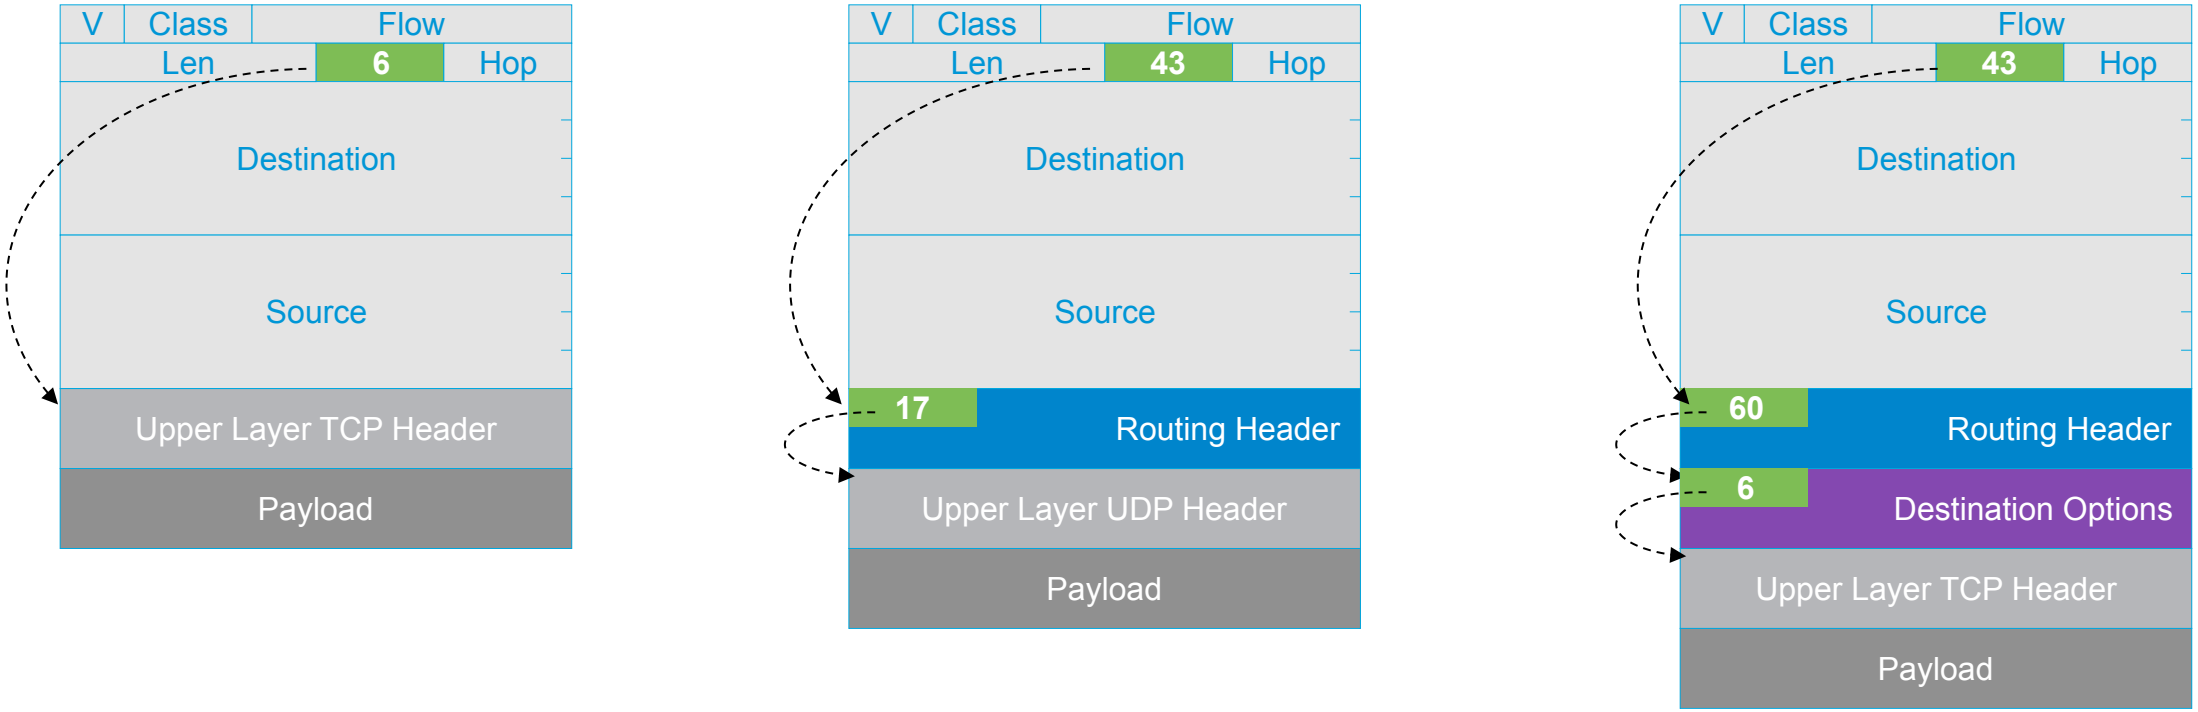

- Extension Headers Are Daisy Chained

© 2011 Cisco and/or its affiliates. All rights reserved.

IPv6 extension headers

Order is important

RFC 2460

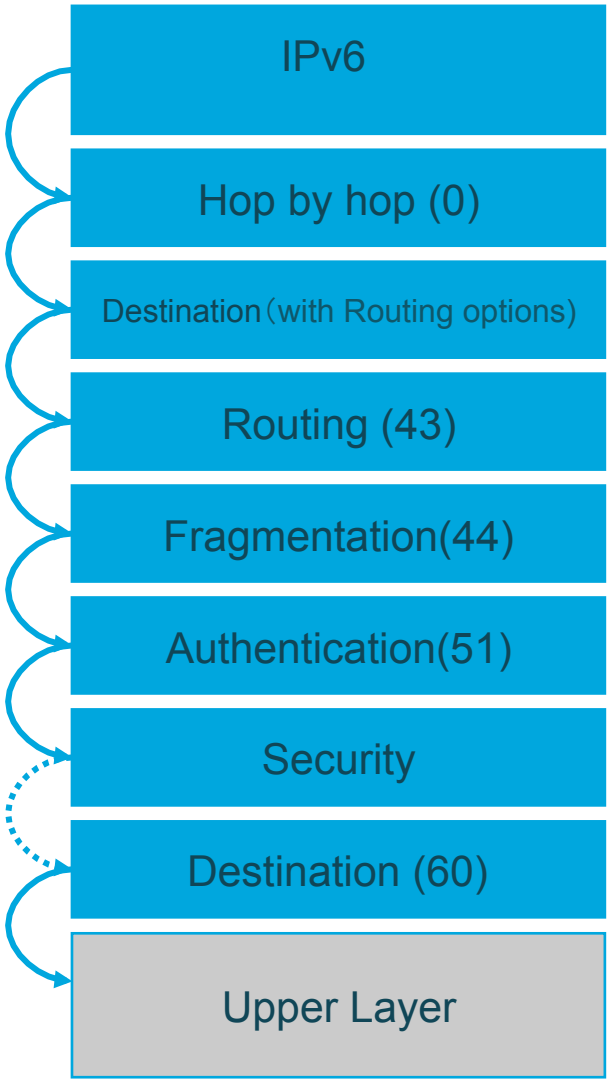

- Hop by hop (0): Processed by every router
- Destination (with Routing options): Processed by routers listed in Routing extension
- Routing (43): List of routers to cross
- Fragmentation(44): Processed by the destination
- Authentication(51): After reassembling the packet
- Security: Cipher the content of the remaining information
- Destination (60): Processed **only** by the destination

© 2011 Cisco and/or its affiliates. All rights reserved.

IPv4 Protocol Stack

The relevant bits

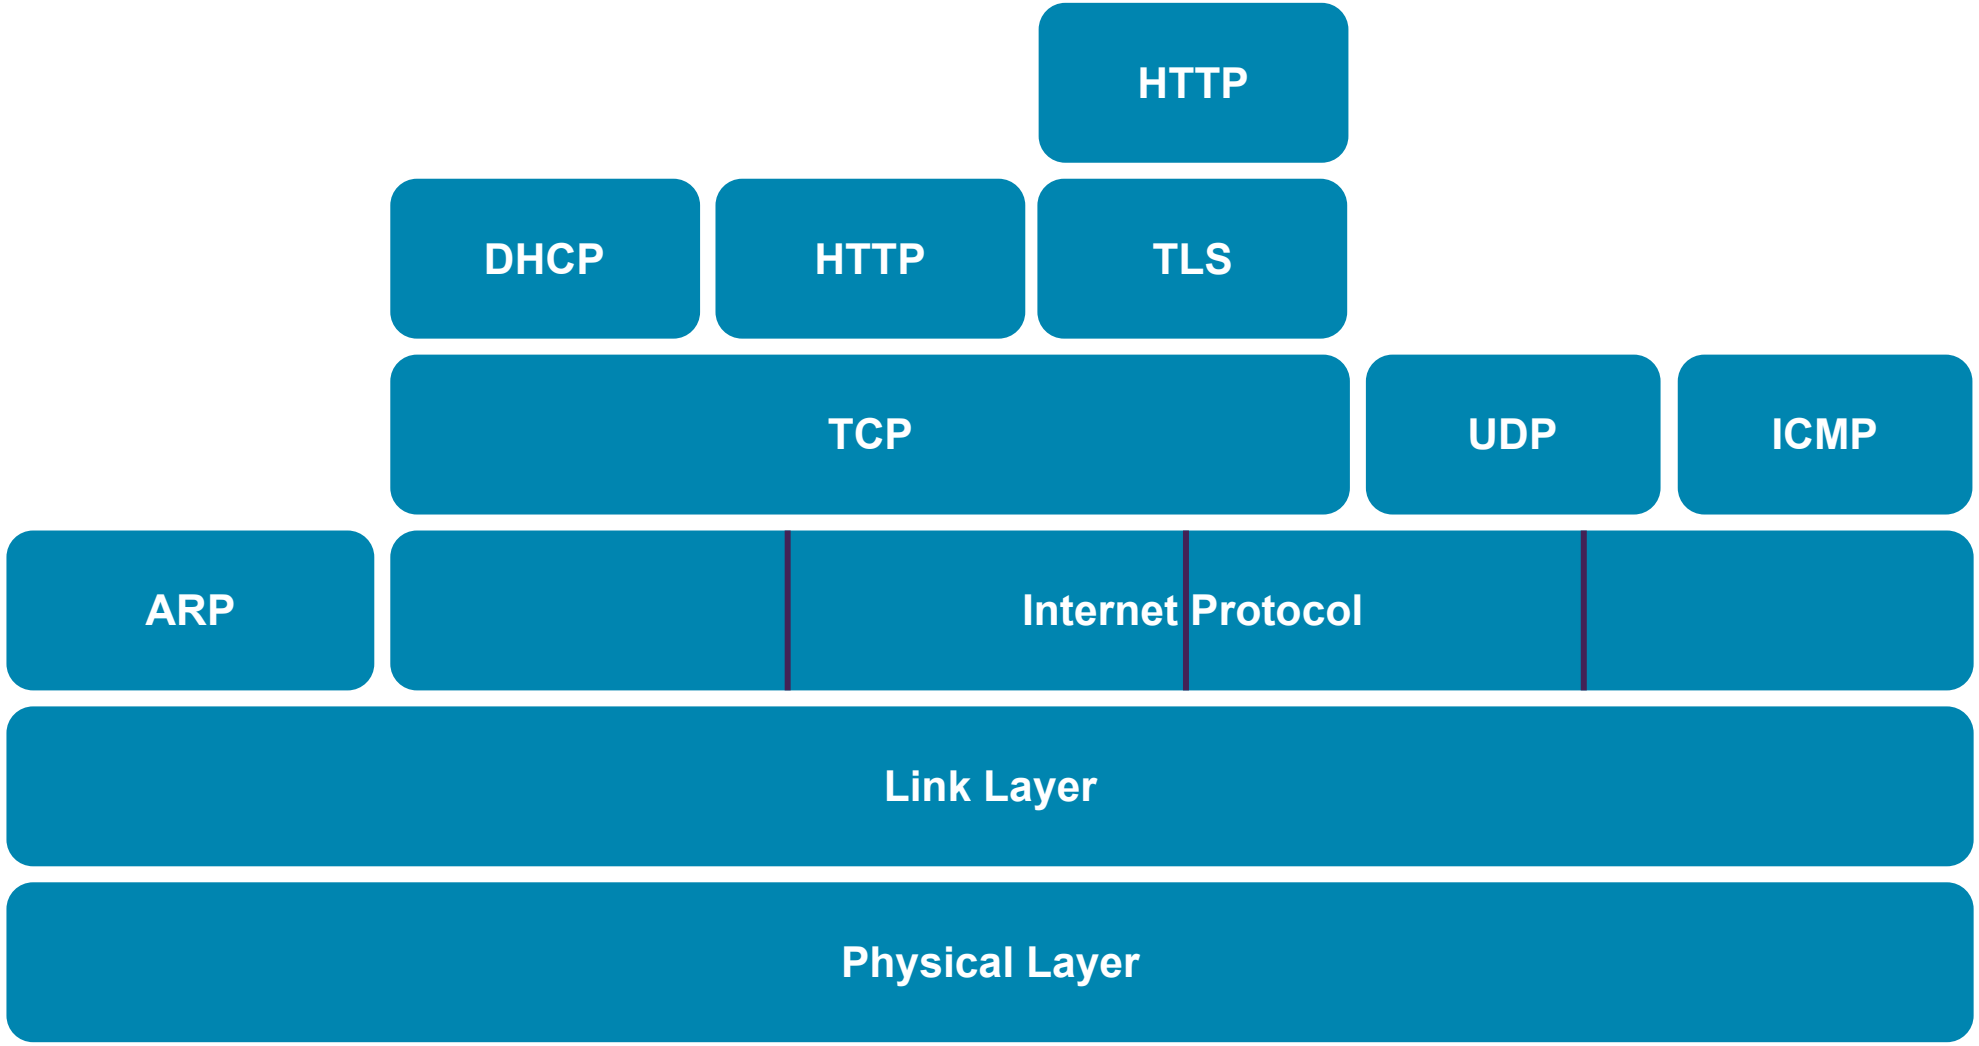

© 2011 Cisco and/or its affiliates. All rights reserved.

IPv6 Protocol Stack

More than just 128 bits

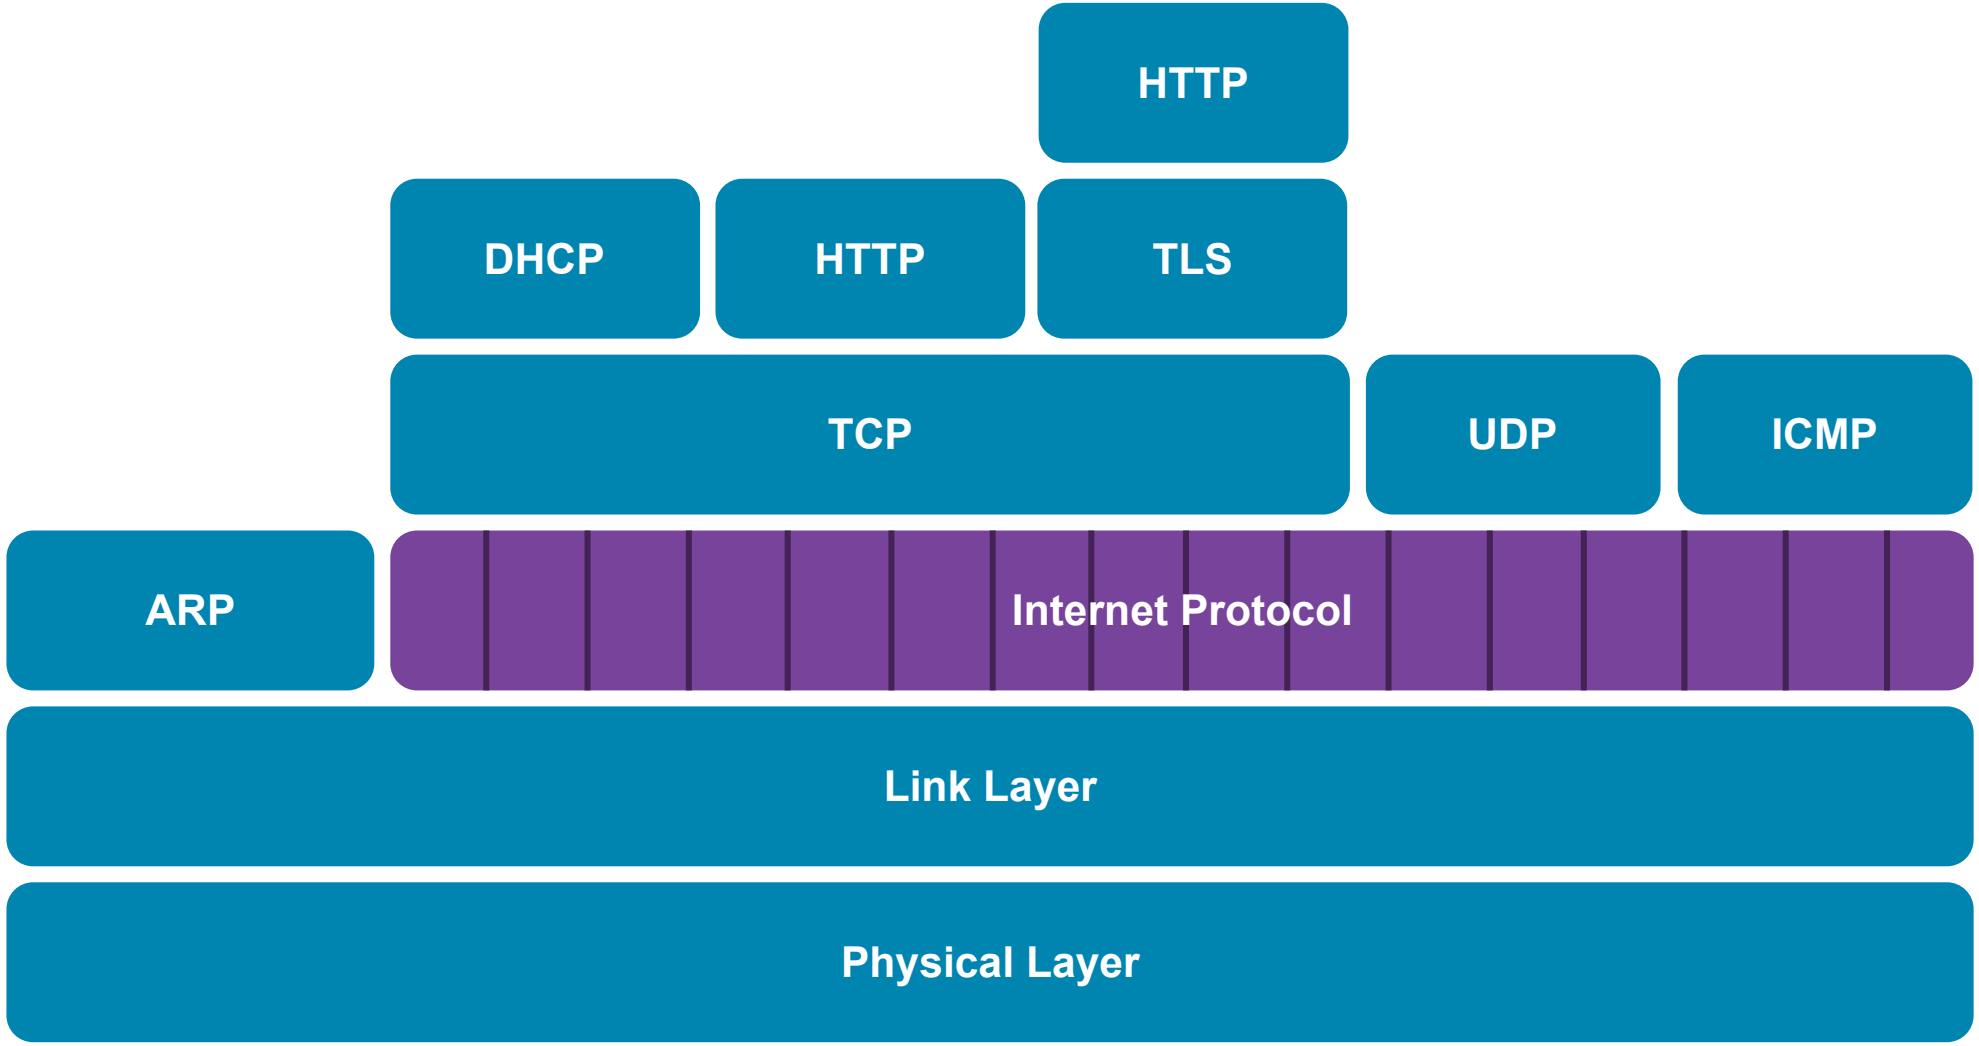

© 2011 Cisco and/or its affiliates. All rights reserved.

IPv6 Protocol Stack

New features

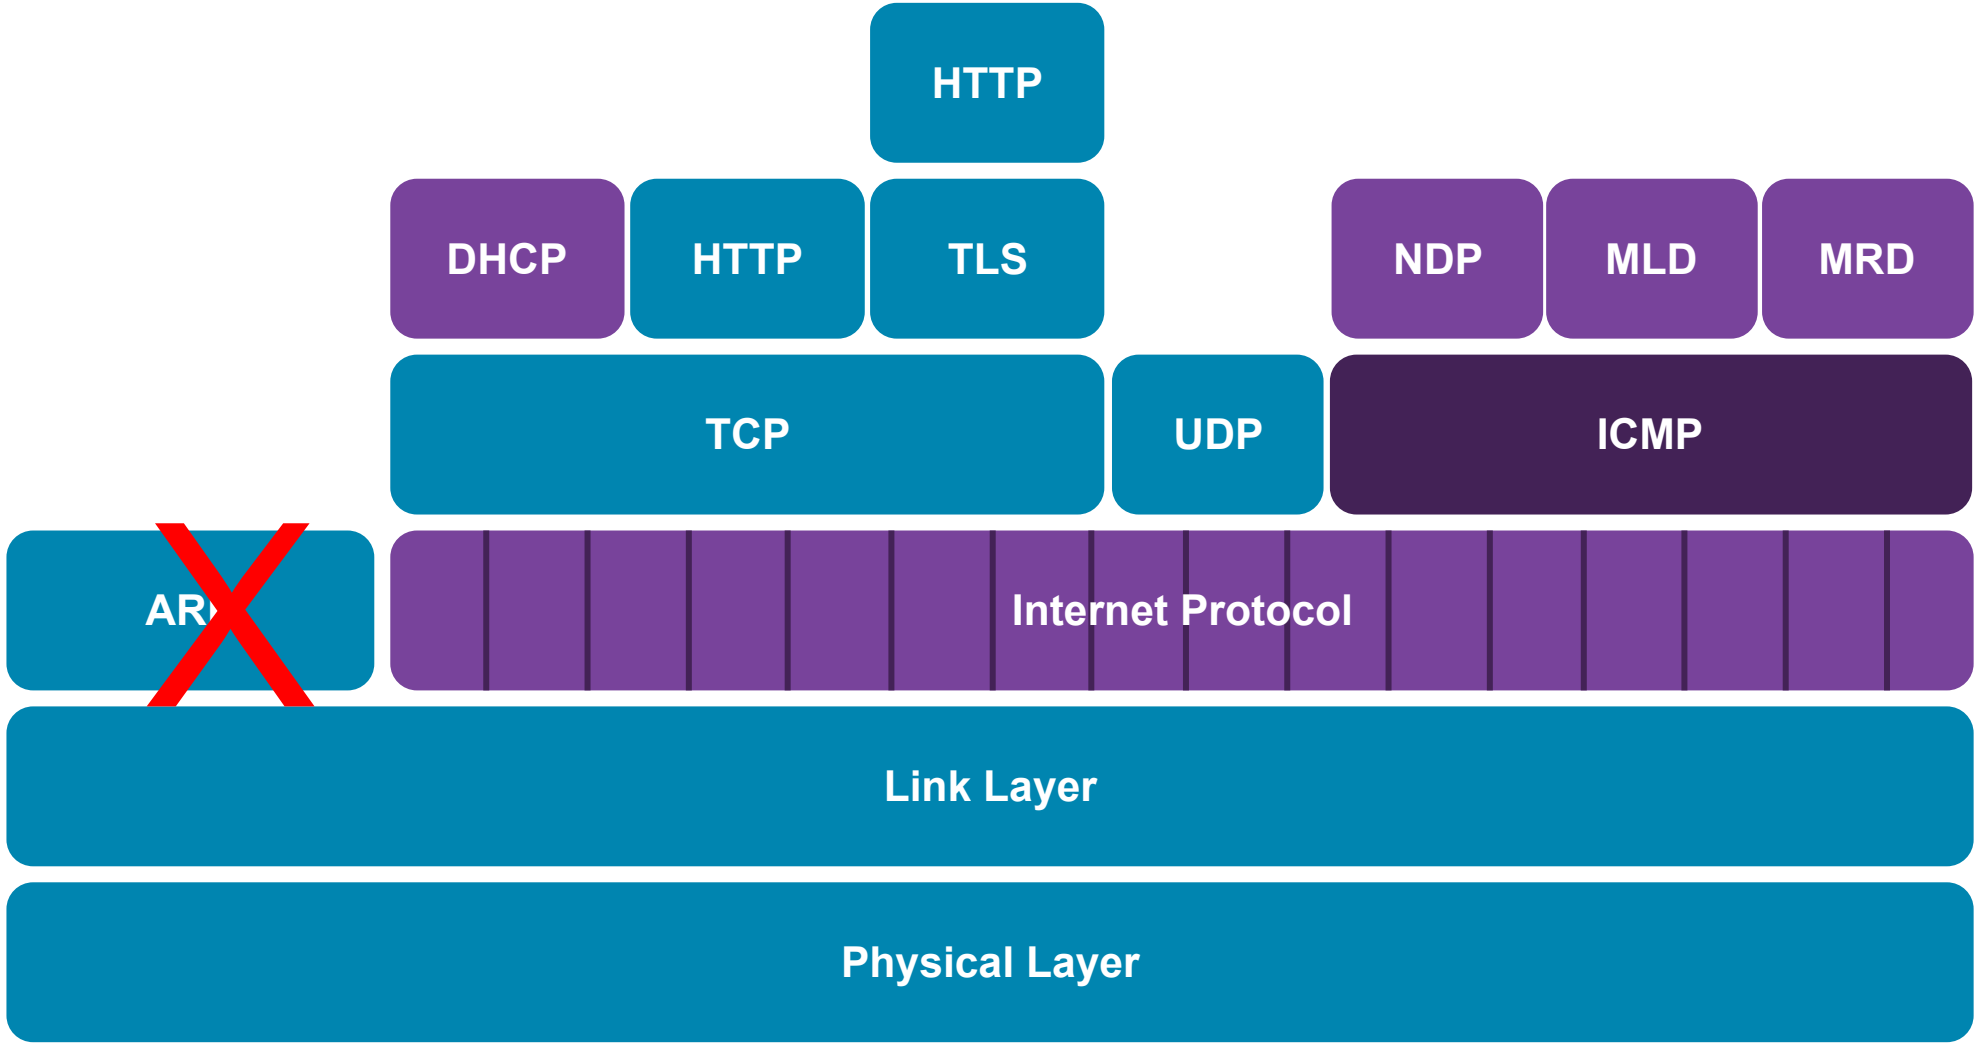

© 2011 Cisco and/or its affiliates. All rights reserved.

IPv4/IPv6 Technology Comparison

Service	IPv4	IPv6
Addressing Range	32-bit, NAT	128-bit, Multiple Scopes
IP Provisioning	Manual, DHCP	Manual, SLAAC, DHCP (and Renumbering capability)
Security	IPSec	IPSec
Mobility	Mobile IP	Mobile IP with Direct Routing
Quality-of-Service	Differentiated Service, Integrated Service	Differentiated Service, Integrated Service
Multicast	IGMP/PIM/MBGP	MLD/PIM/MBGP, Scope Identifier

© 2011 Cisco and/or its affiliates. All rights reserved.

IPv6 Addressing

The Network Portion

© 2011 Cisco and/or its affiliates. All rights reserved.

IPv6 Addresses

- IPv6 addresses are 128 bits long
  - Segmented into 8 groups of four HEX characters (called HEXtets)
  - Separated by a colon (:)
  - Default behaviour is 64 bits for network ID & 64bits for interface ID
  - Network portion is allocated by Internet registries 2^64 (1.8×10^{19})
  - Still leaves us with ~ 3 billion network prefixes for each person on earth

Global Unicast Identifier Example

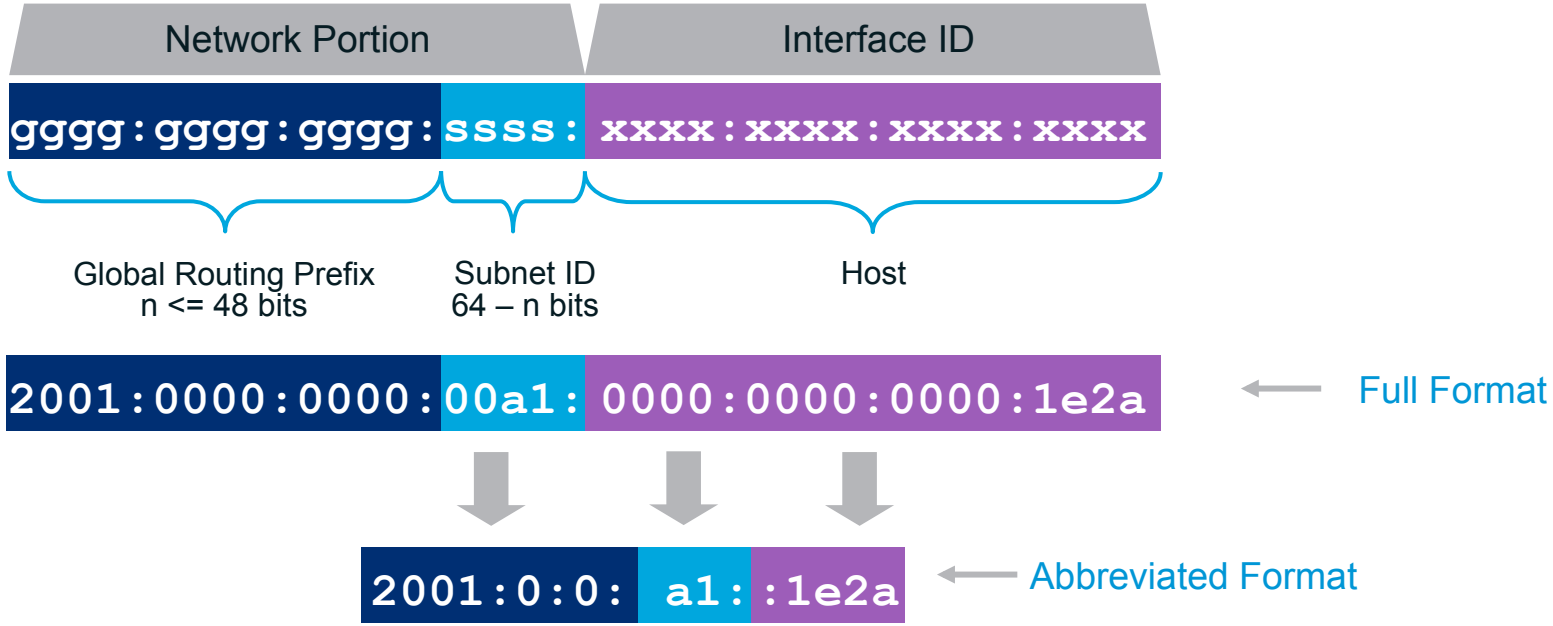

© 2011 Cisco and/or its affiliates. All rights reserved.

IPv6 Address Syntax

- Hex numbers are not case sensitive

  `2001:0db8:0000:130F:0000:0000:087C:AaAa`

  RFC5952 recommends using ‘lowercase’ characters

- Abbreviations are possible

  Zeros in contiguous blocks can be represented by (::)

  `2001:0db8:0000:130f:0000:0000:087c:aaaa`

  `2001:0db8:0000:130f::87c:aaaa`

  Double colon can only appear once in the address

  RFC5952 recommends applying the double colon to the longest all zeros block

- Leading zeros can be omitted

  `2001:0db8:0000:130f::87c:aaaa`

  `2001:0db8:0:130f::87c:aaaa`

- Trailing zeros cannot be omitted

  `2001:0db8:0012::/48 = 2001:db8:12::/48`

  `2001:db80:1200::/48 ≠ 2001:db8:12::/48`

© 2011 Cisco and/or its affiliates. All rights reserved.

IPv6 Address Syntax

- IPv6 uses CIDR representation
  - IPv4 address looks like `98.10.0.0/16`
  - IPv6 address is represented the same way `2001:db8:12::/48`
- Loopback address representation
  - `0:0:0:0:0:0:0:1 == ::1`
  - Same as 127.0.0.1 in IPv4, it identifies self
- Unspecified address representation
  - `0:0:0:0:0:0:0:0 == ::`
  - Used as a placeholder when no address available
  - (Initial DHCP request, Duplicate Address Detection DAD)
  - NOT the default route
- Default Route representation
  - `::/0`

© 2011 Cisco and/or its affiliates. All rights reserved.

IPv6 Address Scopes

- Addresses are assigned to interfaces
- An IPv6 interface is “expected” to have multiple addresses and multiple scopes
- Addresses have scope
  - Link Local
  - Site Local
    - Deprecated (RFC3879). Replaced by Unique Local which is part of the Global scope
  - Global
- Addresses have lifetime
  - Valid and preferred lifetime
- All scopes utilise Multicast

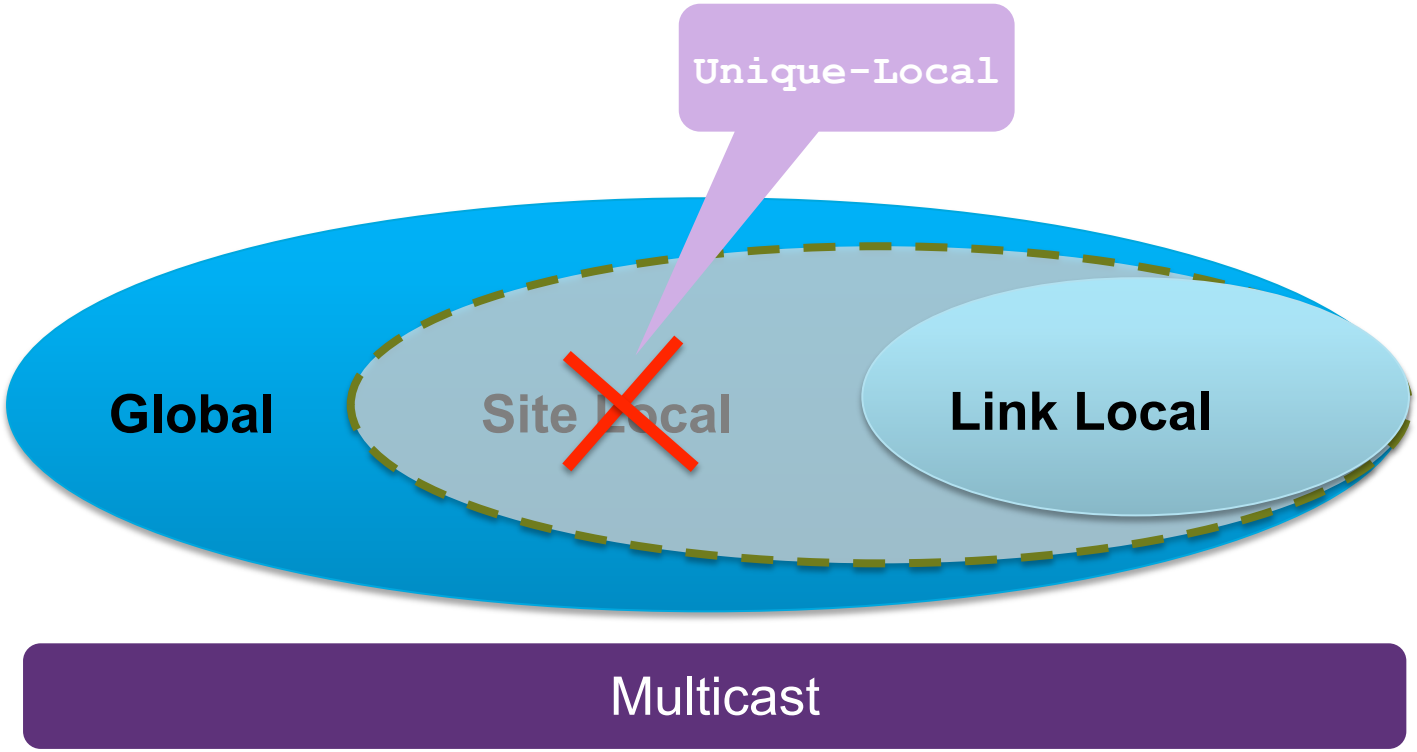

© 2011 Cisco and/or its affiliates. All rights reserved.

IPv6 Address Types

- Three types of unicast address scopes

  Link-Local – Non routable exists on single layer 2 domain **(FE80::/64)**

FE80:0000:0000:0000:	xxxx:xxxx:xxxx:xxxx

  Unique-Local – Routable within administrative domain **(FC00::/7)**

FCgg:gggg:gggg:	ssss:	xxxx:xxxx:xxxx:xxxx
FDgg:gggg:gggg:	ssss:	xxxx:xxxx:xxxx:xxxx

  Global – Routable across the Internet **(2000::/3)**

2ggg:gggg:gggg:	ssss:	xxxx:xxxx:xxxx:xxxx
3ggg:gggg:gggg:	ssss:	xxxx:xxxx:xxxx:xxxx

- Multicast addresses **(FF00::/8)**

FFzs:	xxxx:xxxx:xxxx:xxxx:xxxx:xxxx:xxxx

  Flags (z) in 3rd nibble (4 bits) Scope (s) into 4th nibble

© 2011 Cisco and/or its affiliates. All rights reserved.

Link-Local Address

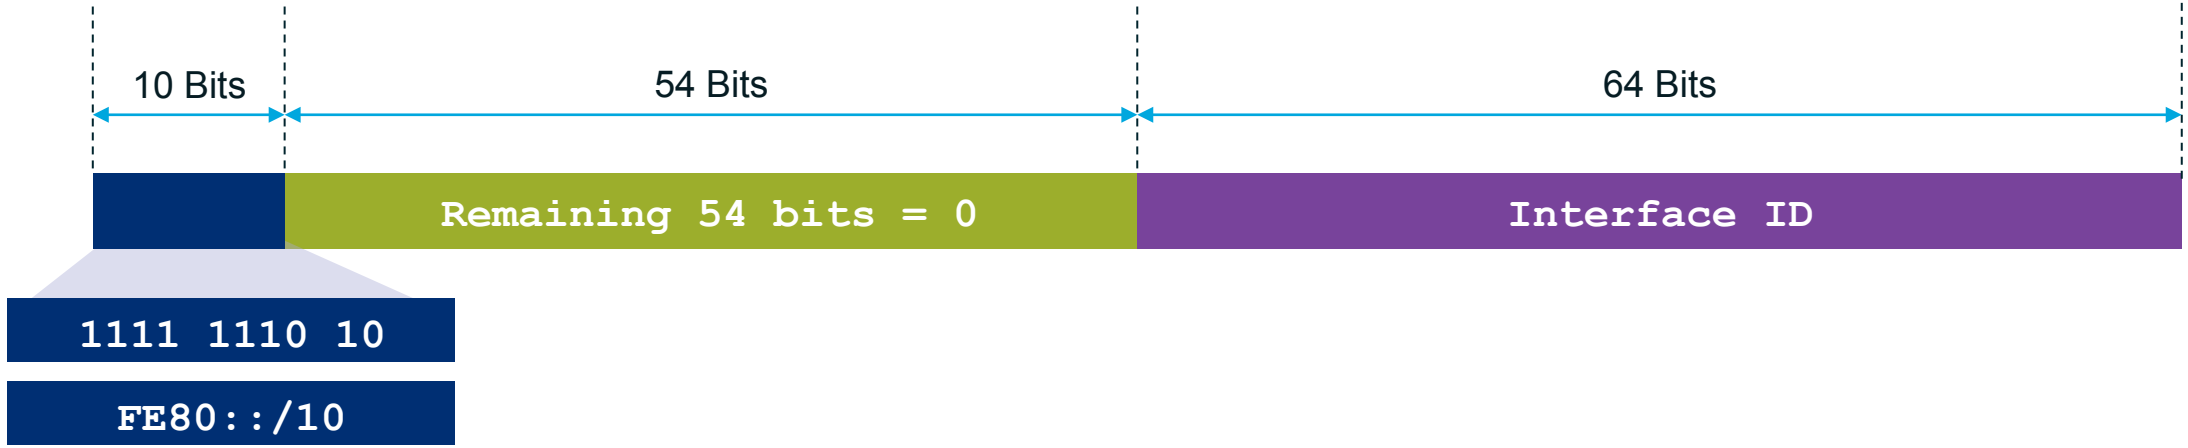

- Mandatory for communication between two IPv6 devices on the same segment
- Automatically self assigned by the device using EUI-64
- Also used for next-hop calculation in routing protocols
- Only link specific scope
- Remaining 54 bits could be zero or any manually configured value

© 2011 Cisco and/or its affiliates. All rights reserved.

Link local address

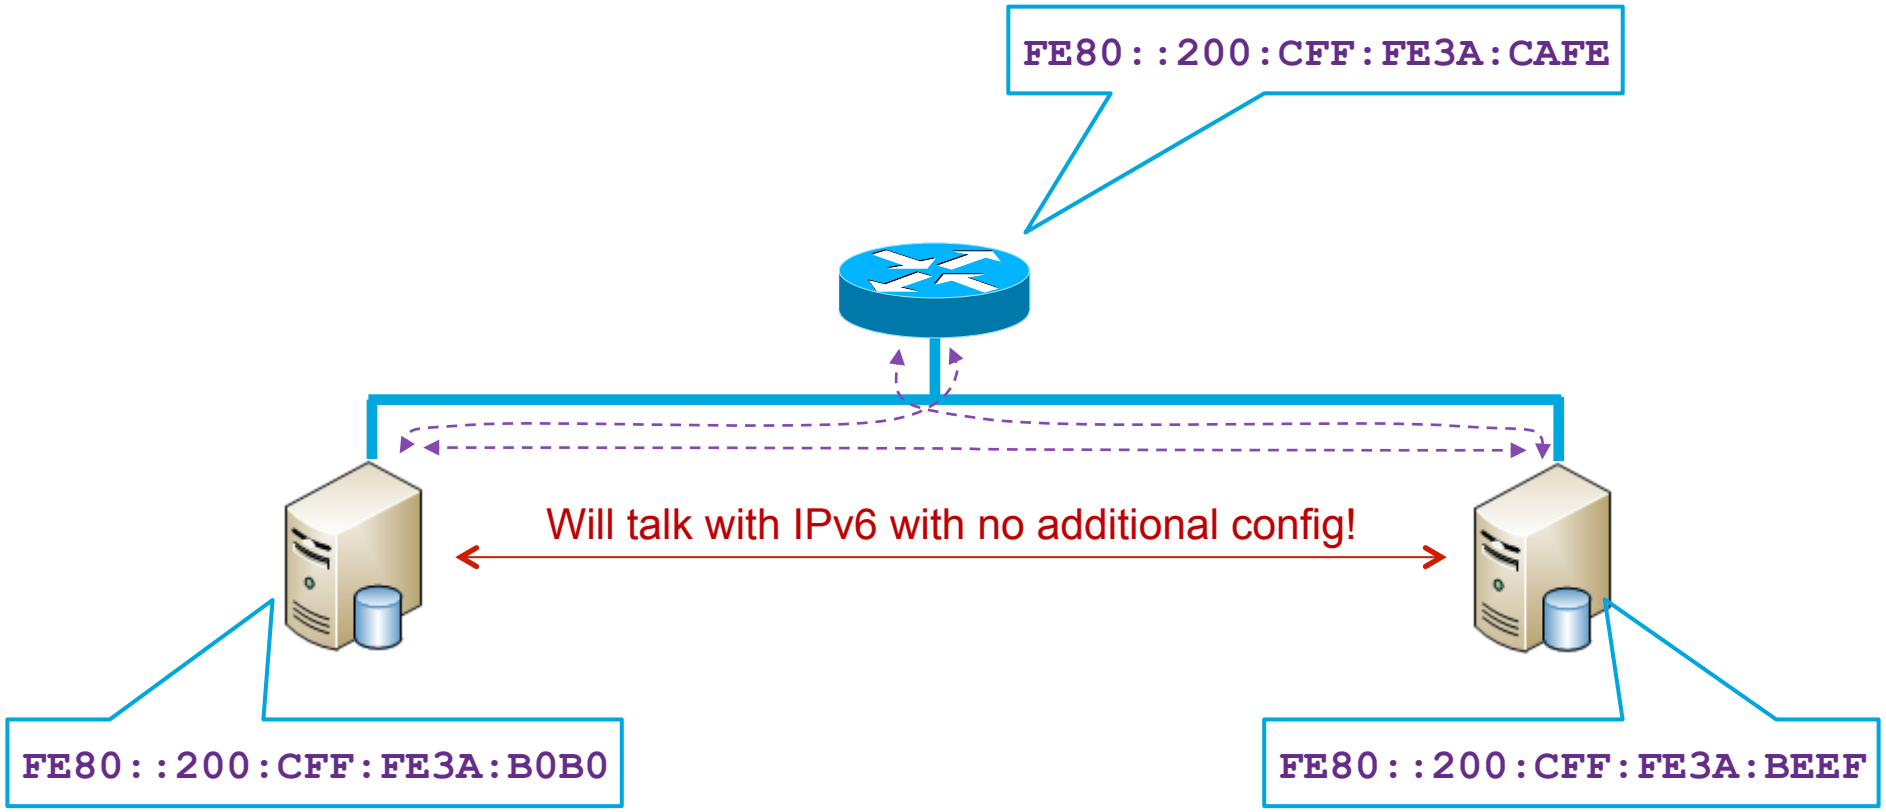

- Once devices are IPv6 enabled and have link local address they will talk over IPv6.
- Beware, you may already be running IPv6 and not know it !

© 2011 Cisco and/or its affiliates. All rights reserved.

Unique Local Address (RFC 4193)

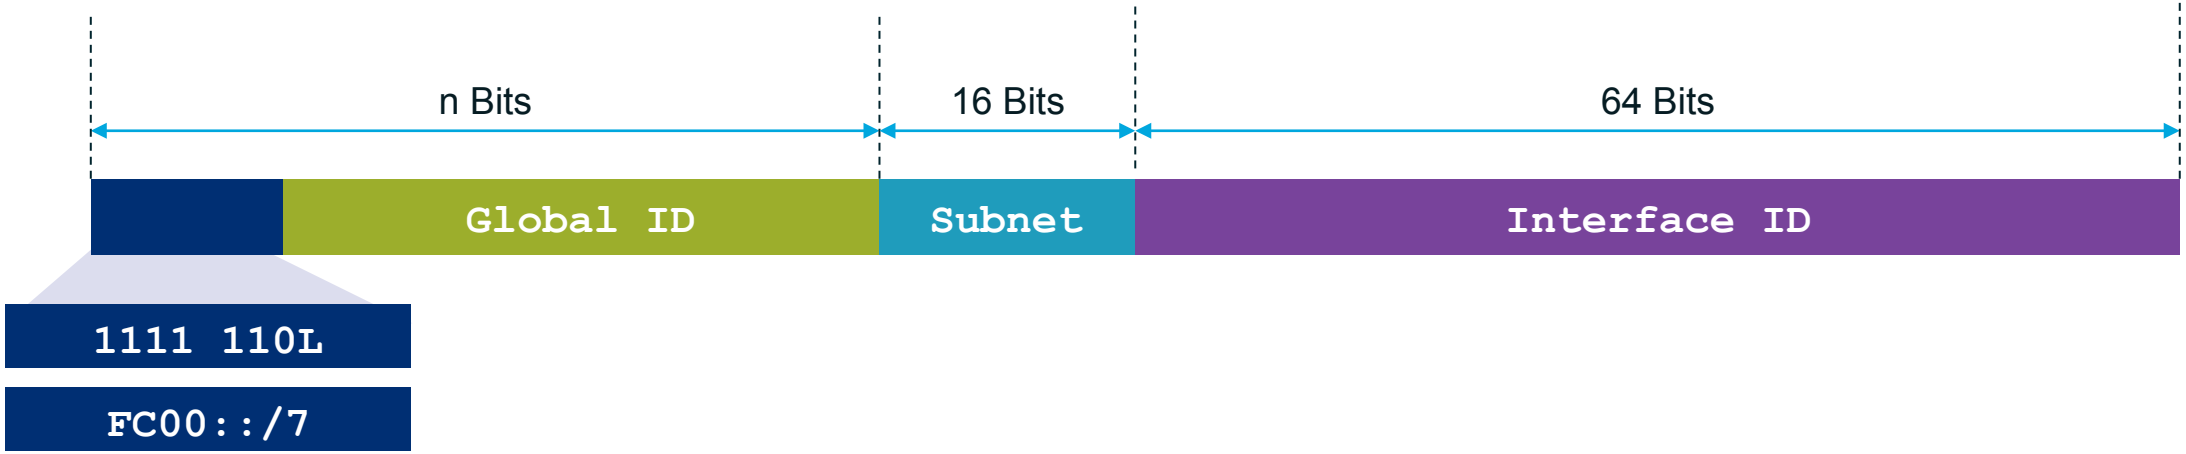

- ULA are “like” RFC 1918 – not routable on Internet through routing policy
- ULA uses include
  - Local communications
  - Inter-site VPNs (Mergers and Acquisitions)
- FC00::/8 is Registry Assigned (L bit = 0), FD00::/8 is self generated (L bit = 1)
  - Registries not yet assigning ULA space, http://www.sixxs.net/tools/grh/ula/
- Global ID can be generated using an algorithm
  - Low order 40 bits result of SHA-1 Digest {EUI-64 && Time}

© 2011 Cisco and/or its affiliates. All rights reserved.

Global Unicast Addresses

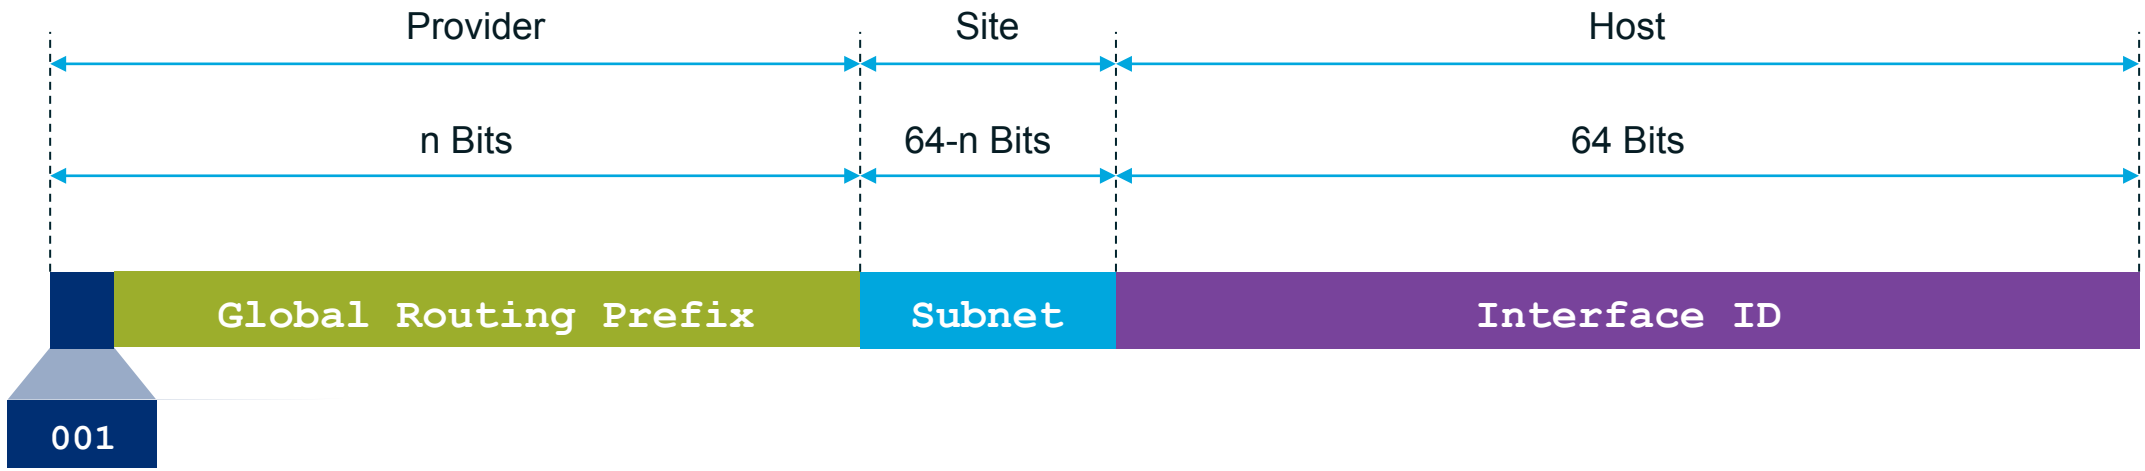

- Addresses for generic use of IPv6
- Globally routable
- Considered best practice for all device numbering
- Common allocation sizes are /32, /48, /52, /56, /64
- Consider using 2000::/3 as the default route

© 2011 Cisco and/or its affiliates. All rights reserved.

IPv6 Addressing

The Interface ID

© 2011 Cisco and/or its affiliates. All rights reserved.

The IPv6 Address Interface ID

- Interface ID of unicast address may be assigned in different ways
  - Auto-configured from a 64-bit EUI-64 or expanded from a 48-bit MAC
  - Auto-generated pseudo-random number (to address privacy concerns)
  - Assigned via DHCP
  - Manually configured
- EUI-64 format to do stateless auto-configuration
  - Expands the 48 bit MAC address to 64 bits by inserting FFFE into the middle
  - To ensure chosen address is from a unique Ethernet MAC address
  - The universal/local (“u” bit) is set to 1 for global scope and 0 for local scope

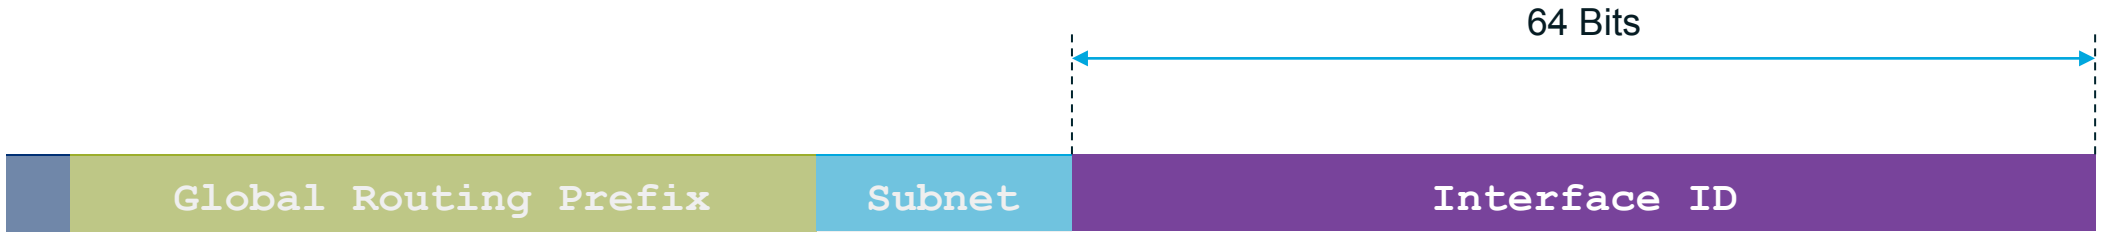

© 2011 Cisco and/or its affiliates. All rights reserved.

IPv6 Interface Identifier (EUI-64 format)

- Cisco uses the EUI-64 format to do stateless auto-configuration
- This format expands the 48 bit MAC address to 64 bits by inserting FFFE into the middle 16 bits
- To make sure that the chosen address is from a unique Ethernet MAC address, the universal/local (“u” bit) is set to 1 for global scope and 0 for local scope

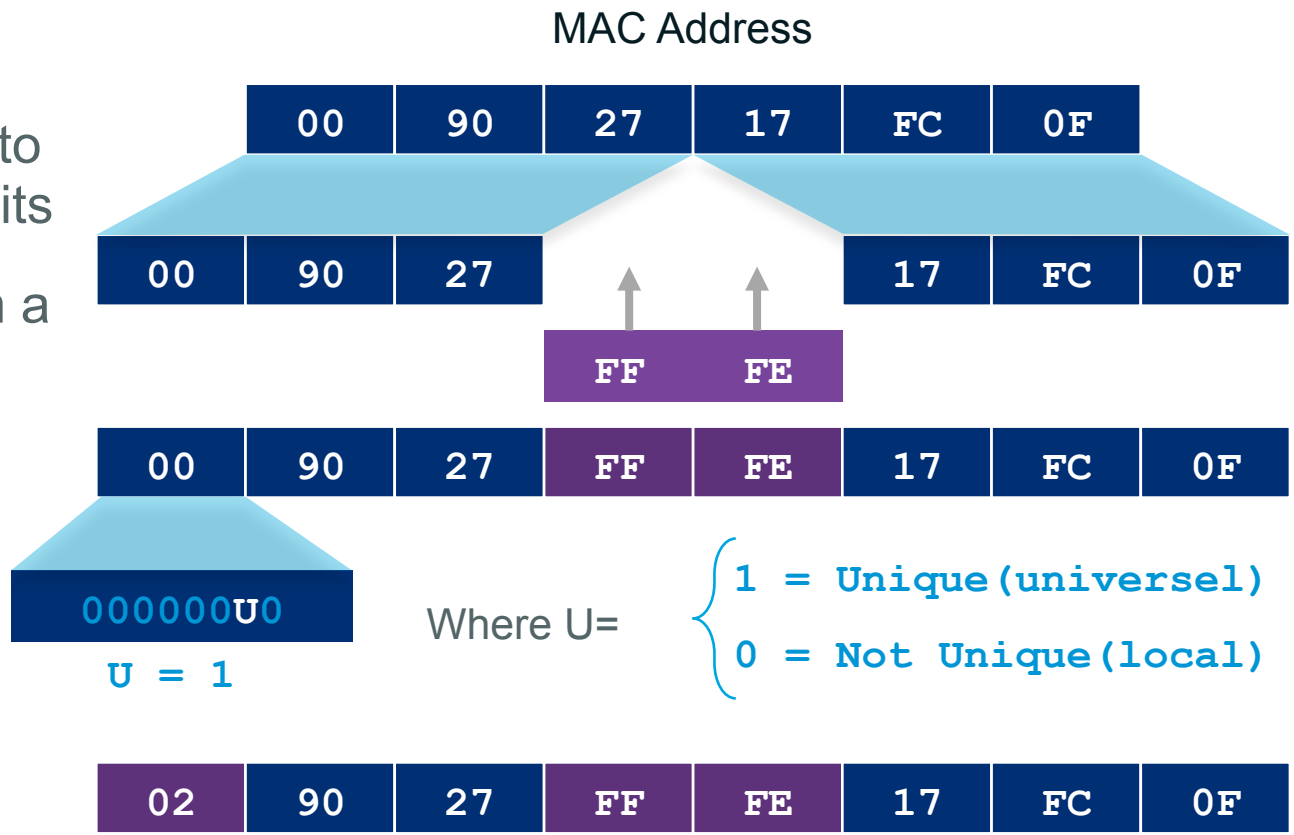

© 2011 Cisco and/or its affiliates. All rights reserved.

Fragmentation and Path MTU Discovery

© 2011 Cisco and/or its affiliates. All rights reserved.

Fragmentation in IPv6

NOTE: Fragmentation can _only_ be done by IPv6 hosts. Intermediary devices can not fragment !

- The original large packet consists of two parts
- Unfragmentable part
  - IPv6 header plus any headers that must be processed by the nodes en-route
  - Unfragmentable part is repeated with fragments appended to it following the “fragment header”
- Fragmentable part
  - The headers that need to be processed only by the destination node = the end-to-end headers + upper layer header and data
  - Fragmentable part is divided into pieces with length multiple of 8 octets
- RFC 2460 Section 4.5 defines the fragmentation header
- Minimum MTU for IPv6 is 1280 bytes
  - All links MUST support it

© 2011 Cisco and/or its affiliates. All rights reserved.

Fragment Header

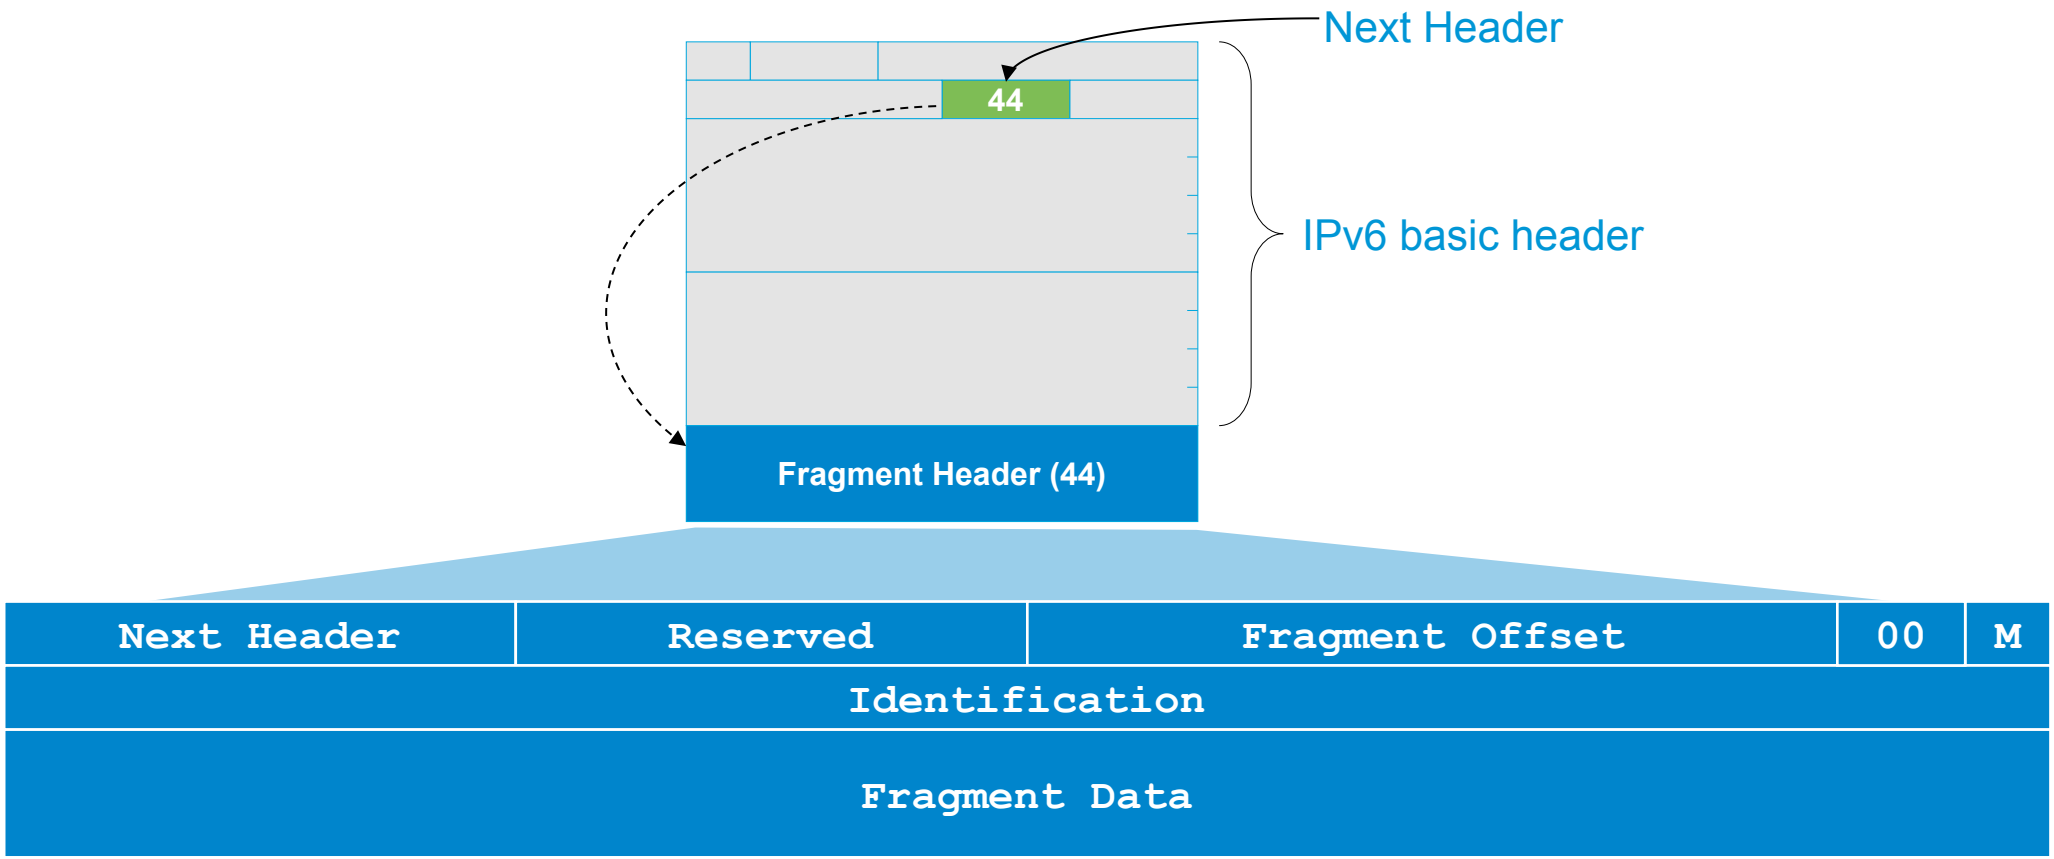

- Fragmentation is left to end devices in IPv6
  - Routers do not perform fragmentation
- Fragment header used when an end node has to send a packet larger than the path MTU

© 2011 Cisco and/or its affiliates. All rights reserved.

Path MTU Discovery

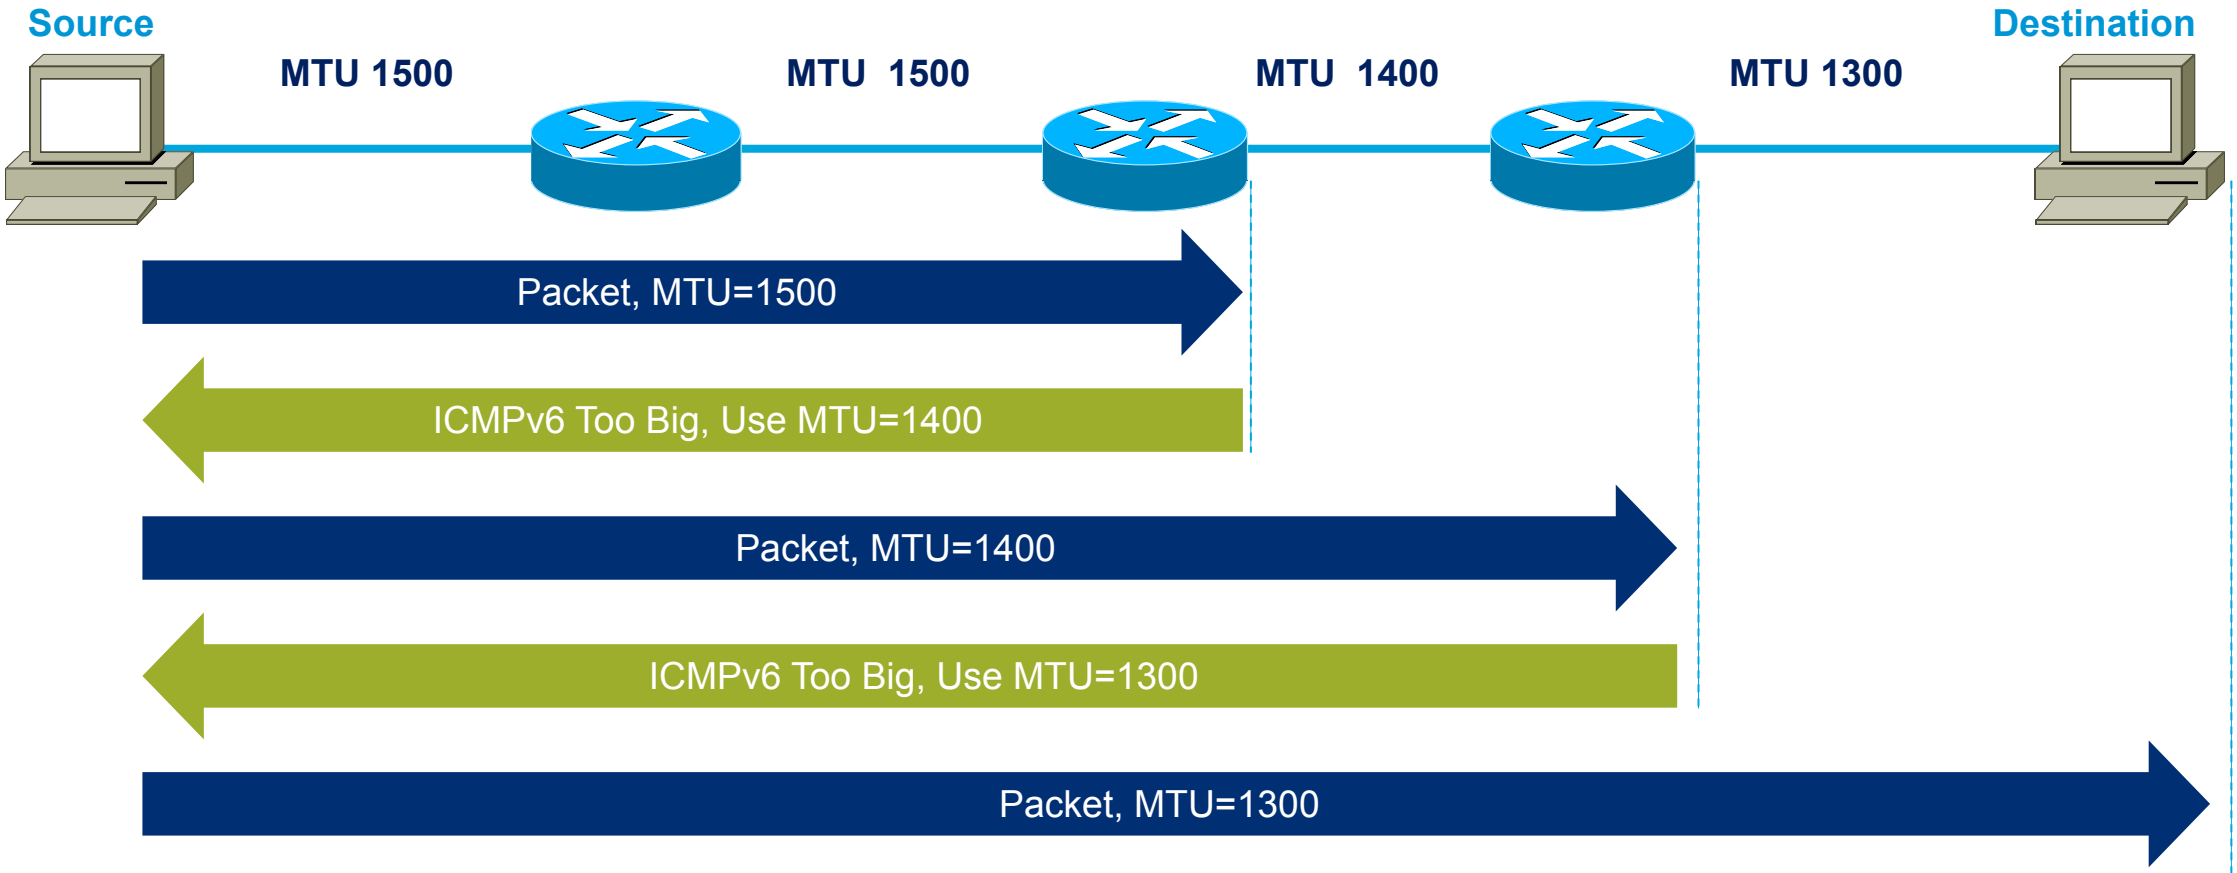

- Store PMTU per destination (if received)
- Age out PMTU (10 mins), reset to first link MTU

© 2011 Cisco and/or its affiliates. All rights reserved.

ICMPv6 and Neighbor Discovery

© 2011 Cisco and/or its affiliates. All rights reserved.

Something new

There is no more
Broadcast.
There is no more ARP

© 2011 Cisco and/or its affiliates. All rights reserved.

ICMPv6

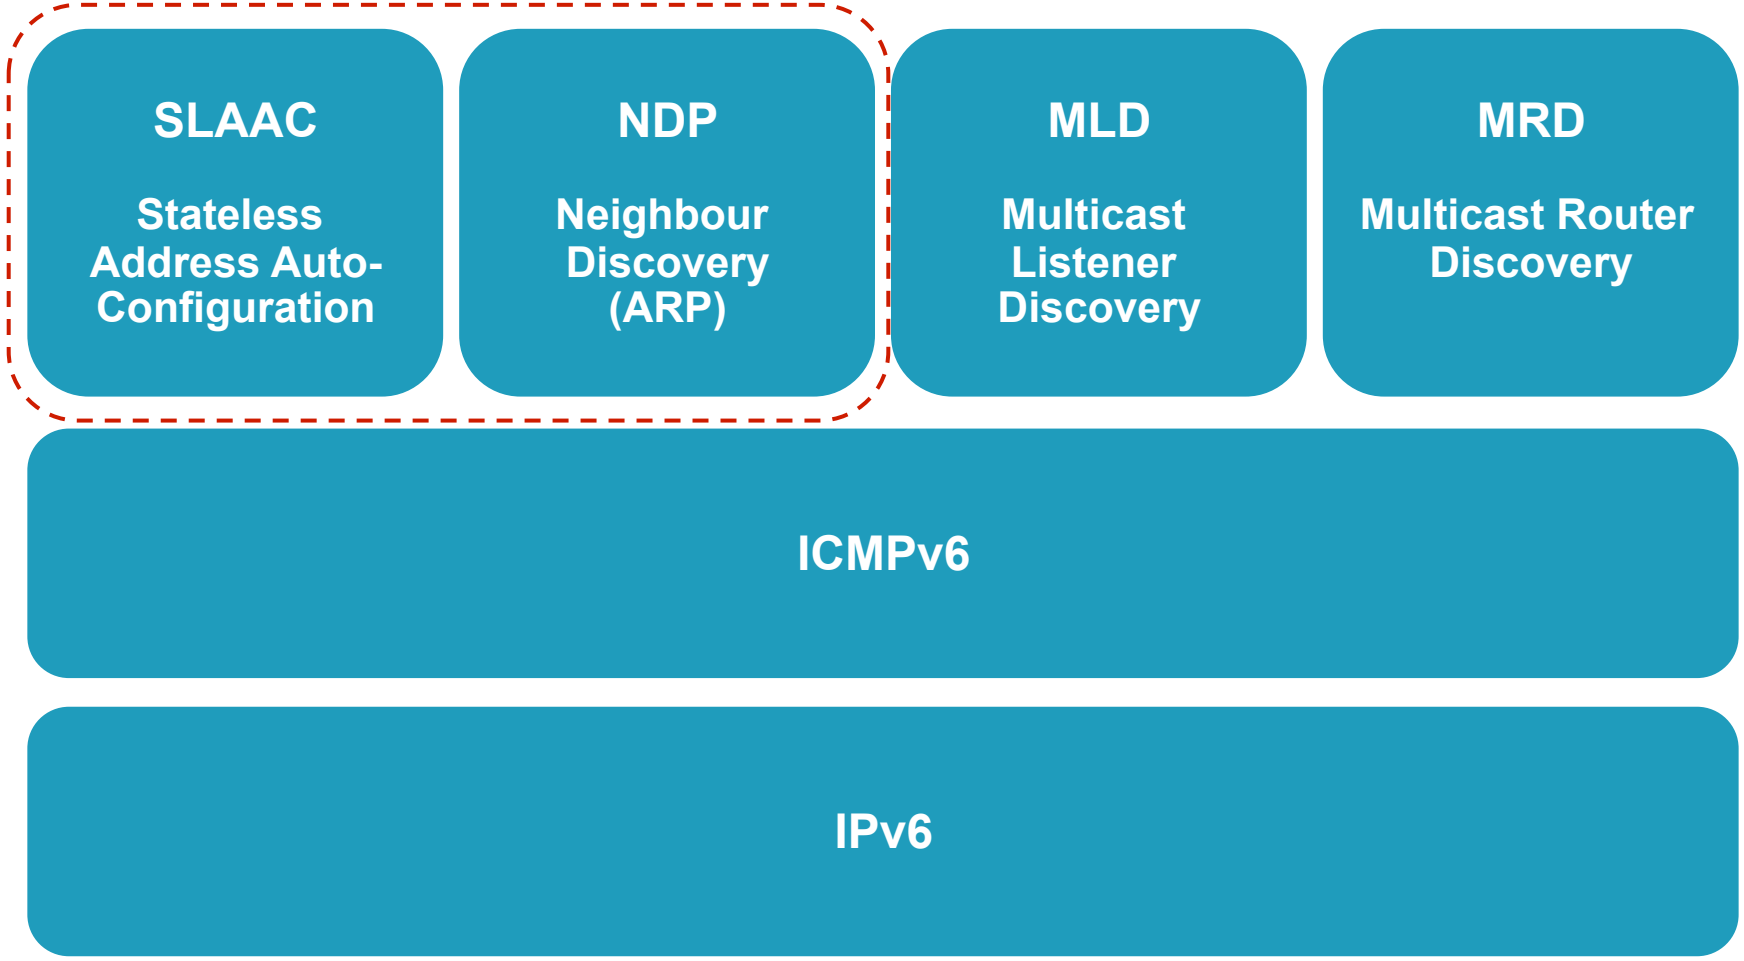

© 2011 Cisco and/or its affiliates. All rights reserved.

ICMPv6 (RFC 2463)

- Internet Control Message Protocol version 6
- Combines several IPv4 functions
  - ICMPv4, IGMP and ARP
- Message types are similar to ICMPv4
  - Destination unreachable (type 1)
  - Packet too big (type 2)
  - Time exceeded (type 3)
  - Parameter problem (type 4)
  - Echo request/reply (type 128 and 129)

© 2011 Cisco and/or its affiliates. All rights reserved.

ICMPv6 Header

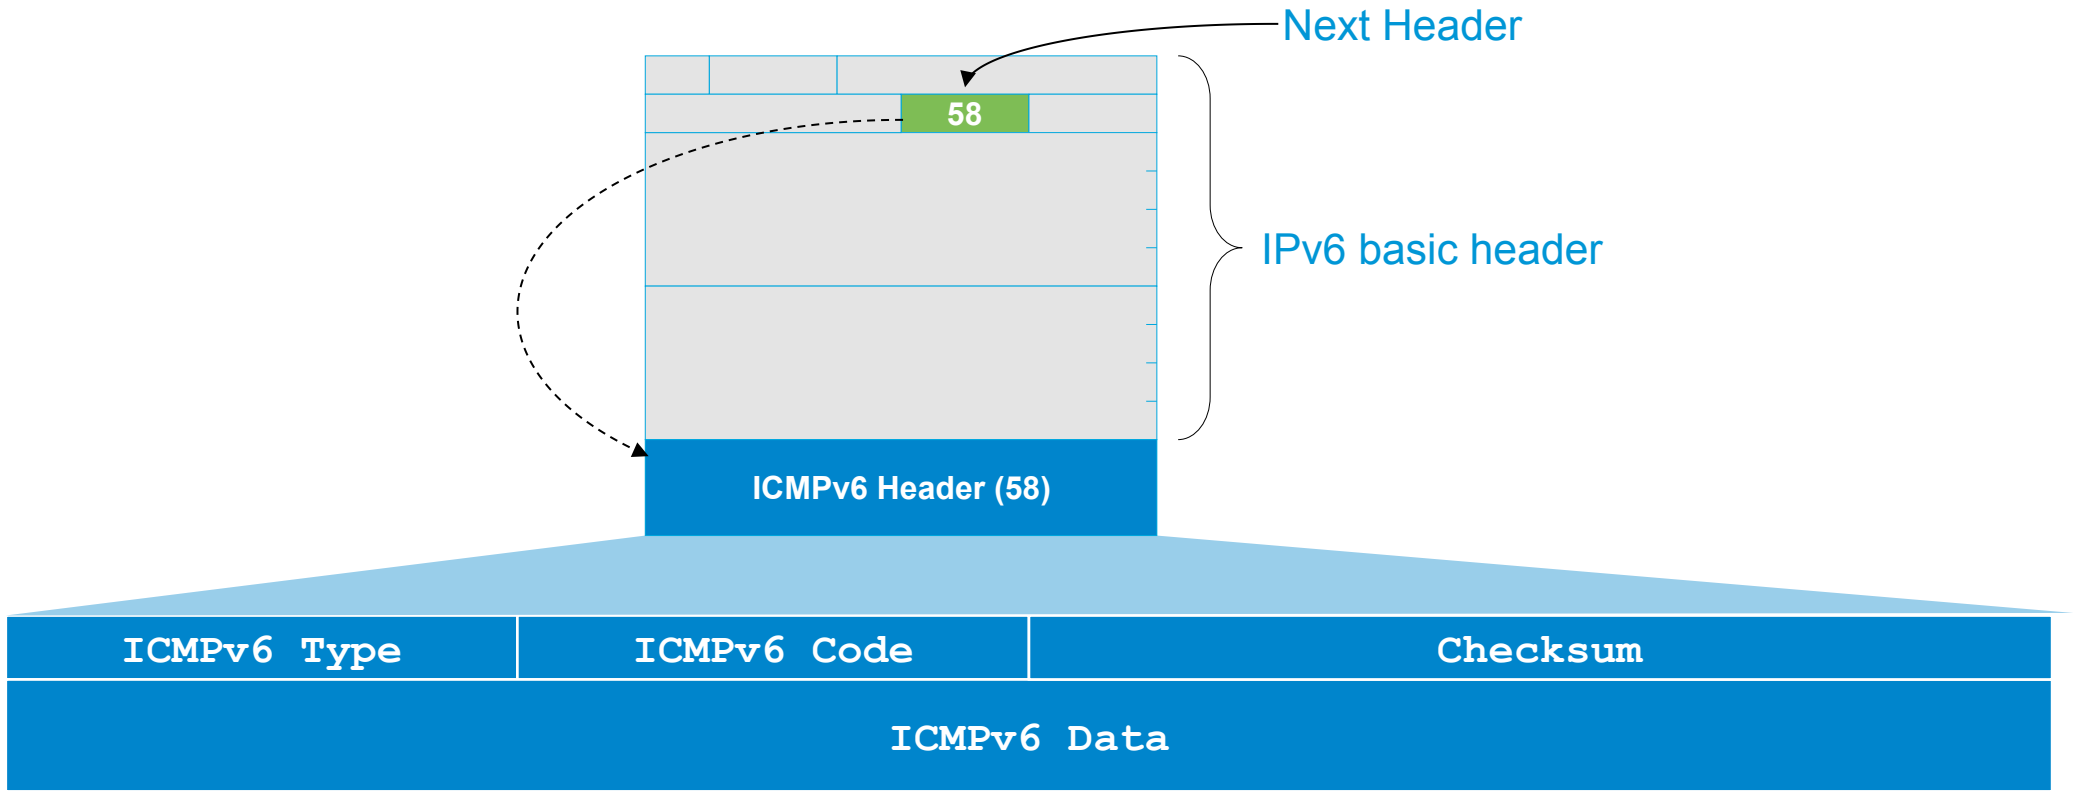

- Also used for Neighbour Discovery, Path MTU discovery and Multicast Listener Discovery (MLD)
  - Type - identifies the message or action needed
  - Code – is a type-specific sub-identifier.
  - Checksum – computed over the entire ICMPv6

© 2011 Cisco and/or its affiliates. All rights reserved.

Neighbor Discovery Messages (ND)

- ND uses ICMPv6 messages
  - Originated from node on link local with a hop limit of 255
  - Receivers checks hop limit is still 255 (has not passed a router)
- Consists of IPv6 header, ICMPv6 header, neighbor discovery header, and neighbor discovery options
- Five neighbor discovery messages

Message	Purpose	ICMP Code	Sender	Target
Router Solicitation (RS)	Prompt routers to send RA	133	Nodes	All routers
Router Advertisement (RA)	Advertise default router, prefixes Operational parameters	134	Routers	Sender of RS All routers
Neighbor Solicitation (NS)	Request link-layer of target	135	Node	Solicited Node Target Node
Neighbor Advertisement (NA)	Response to NS (solicited) Advertise link-layer address change (Unsolicited)	136	Nodes	
Redirect	Inform hosts of a better first hop	137	Routers	

ICMPv6 Neighbor Discovery (RFC 4861)

- Replaces ARP, ICMP (redirects, router discovery)
- Uses ICMPv6 header
- Reachability of neighbours
- Hosts use it to discover routers, auto configuration of addresses (SLAAC)
- Duplicate Address Detection (DAD)

© 2011 Cisco and/or its affiliates. All rights reserved.

IPv4/IPv6 Provisioning Comparison

Function	IPv4	IPv6
Address Assignment	DHCPv4	DHCPv6, SLAAC, Reconfiguration
Address Resolution	ARP RARP	ICMPv6 NS, NA Not Used
Router Discovery	ICMP Router Discovery	ICMPv6 RS, RA
Name Resolution	DNS	DNS

© 2011 Cisco and/or its affiliates. All rights reserved.

Router Solicitation and Advertisement (RS & RA)

Router Solicitation	
ICMP Type	133
IPv6 Source	A Link Local (FE80::1)
IPv6 Destination	All Routers Multicast (FF02::2)
Query	Please send RA

Router Advertisement	
ICMP Type	134
IPv6 Source	A Link Local (FE80::2)
IPv6 Destination	All Nodes Multicast (FF02::1)
Data	Options, subnet prefix, lifetime, autoconfig flag

- Router solicitations (RS) are sent by booting nodes to request RAs for configuring the interfaces
- Routers send periodic Router Advertisements (RA) to the all-nodes multicast address

© 2011 Cisco and/or its affiliates. All rights reserved.

Neighbor Solicitation & Advertisement

- Neighbor Solicitation (NS)

  Used to discover link layer address of IPv6 node

NS Function	Source	Destination
Address resolution	Unicast	Solicited Node Multicast
Node reachability	Unicast	Unicast
Duplicate Address Detection	::0	Solicited Node Multicast

- Neighbor Advertisement (NA)

  Response to neighbor solicitation (NS) message

  A node may also send unsolicited Neighbor Advertisements to announce a link-layer address change.

© 2011 Cisco and/or its affiliates. All rights reserved.

Neighbor Solicitation & Advertisement (NS & NA)

Neighbour Solicitation	
ICMP Type	135
IPv6 Source	A Unicast
IPv6 Destination	B Solicited Node Multicast
Data	FE80:: address of A
Query	What is B link layer address?

Neighbour Advertisment	
ICMP Type	136
IPv6 Source	B Unicast
IPv6 Destination	A Unicast
Data	FE80:: address of B, MAC Address

© 2011 Cisco and/or its affiliates. All rights reserved.

Neighbor Unreachability Detection

- Neighbor is declared reachable if
  - The connection is making forward progress
  - Previously sent data is known to have been delivered correctly
  - Source receives an NA in response to NS
- If neighbor status unknown then send NS
- Defined in RFC 4861 Section 7.3

© 2011 Cisco and/or its affiliates. All rights reserved.

Neighbor Cache Entry States

- INCOMPLETE

  Address resolution is in progress and the link-layer address of the neighbor has not yet been determined

- REACHABLE

  The neighbor is known to have been reachable recently (within tens of seconds ago)

- STALE

  The neighbor is no longer known to be reachable but until traffic is sent to the neighbor, no attempt should be made to verify its reachability

- DELAY

  Delay sending probes for a short while in order to give upper layer protocols a chance to provide reachability confirmation

- PROBE

  The neighbor is no longer known to be reachable, and unicast Neighbor Solicitation probes are being sent to verify reachability

© 2011 Cisco and/or its affiliates. All rights reserved.

Duplicate Address Detection (DAD)
equal to ipv4 gratuitous ARP

Neighbour Solicitation	
ICMP Type	135 (Neighbour Solicitation)
Ethernet DA	33-33-FF-52-F9-D8
IPv6 Header	
IPv6 Source	::
IPv6 Destination	FF02::1:FF52:F9D8
Hop Limit	255
NS Header	
Target Address	FE80::260:8FF:FE52:F9D8

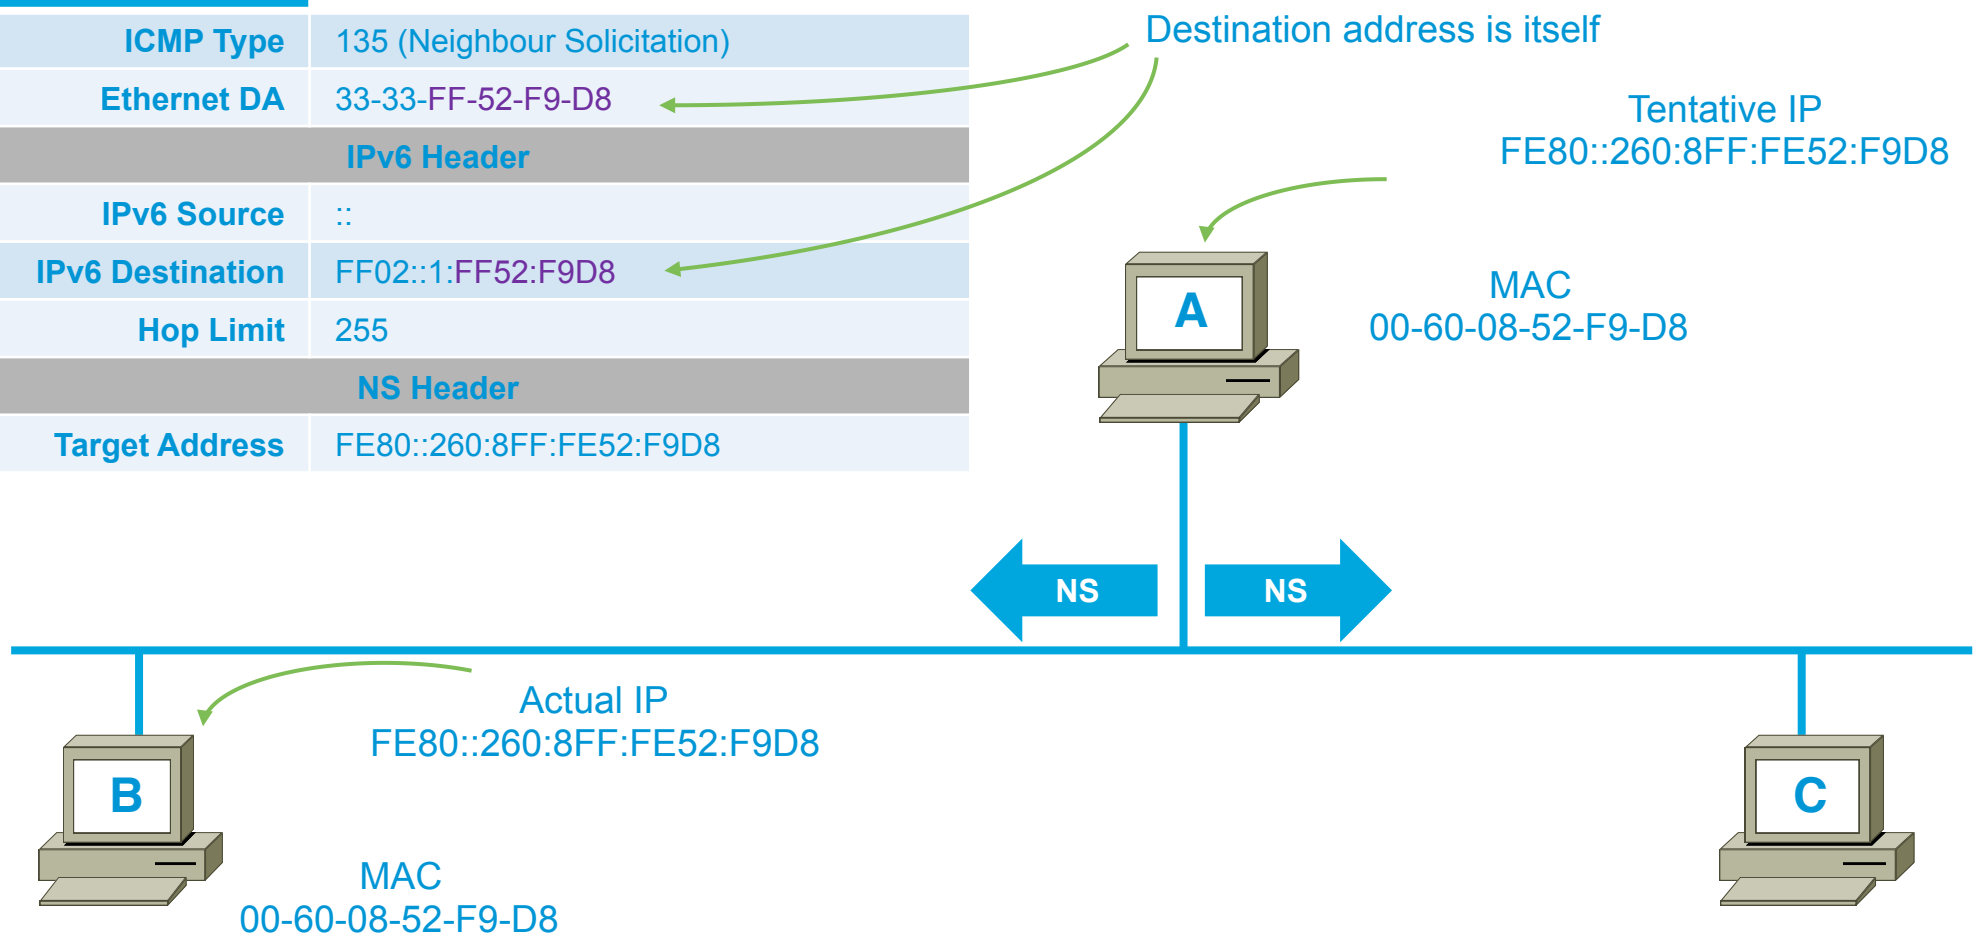

© 2011 Cisco and/or its affiliates. All rights reserved.

Duplicate Address Detection Response

Neighbour Solicitation	
ICMP Type	135 (Neighbour Solicitation)
Ethernet DA	33-33-00-00-00-01
IPv6 Header	
IPv6 Source	FE80::260:8FF:FE52:F9D8
IPv6 Destination	FF02::1
Hop Limit	255
NA Header	
Target Address	FE80::260:8FF:FE52:F9D8
Neighbour Discovery Option	
Target L2 Address	00-60-08-52-F9-D8

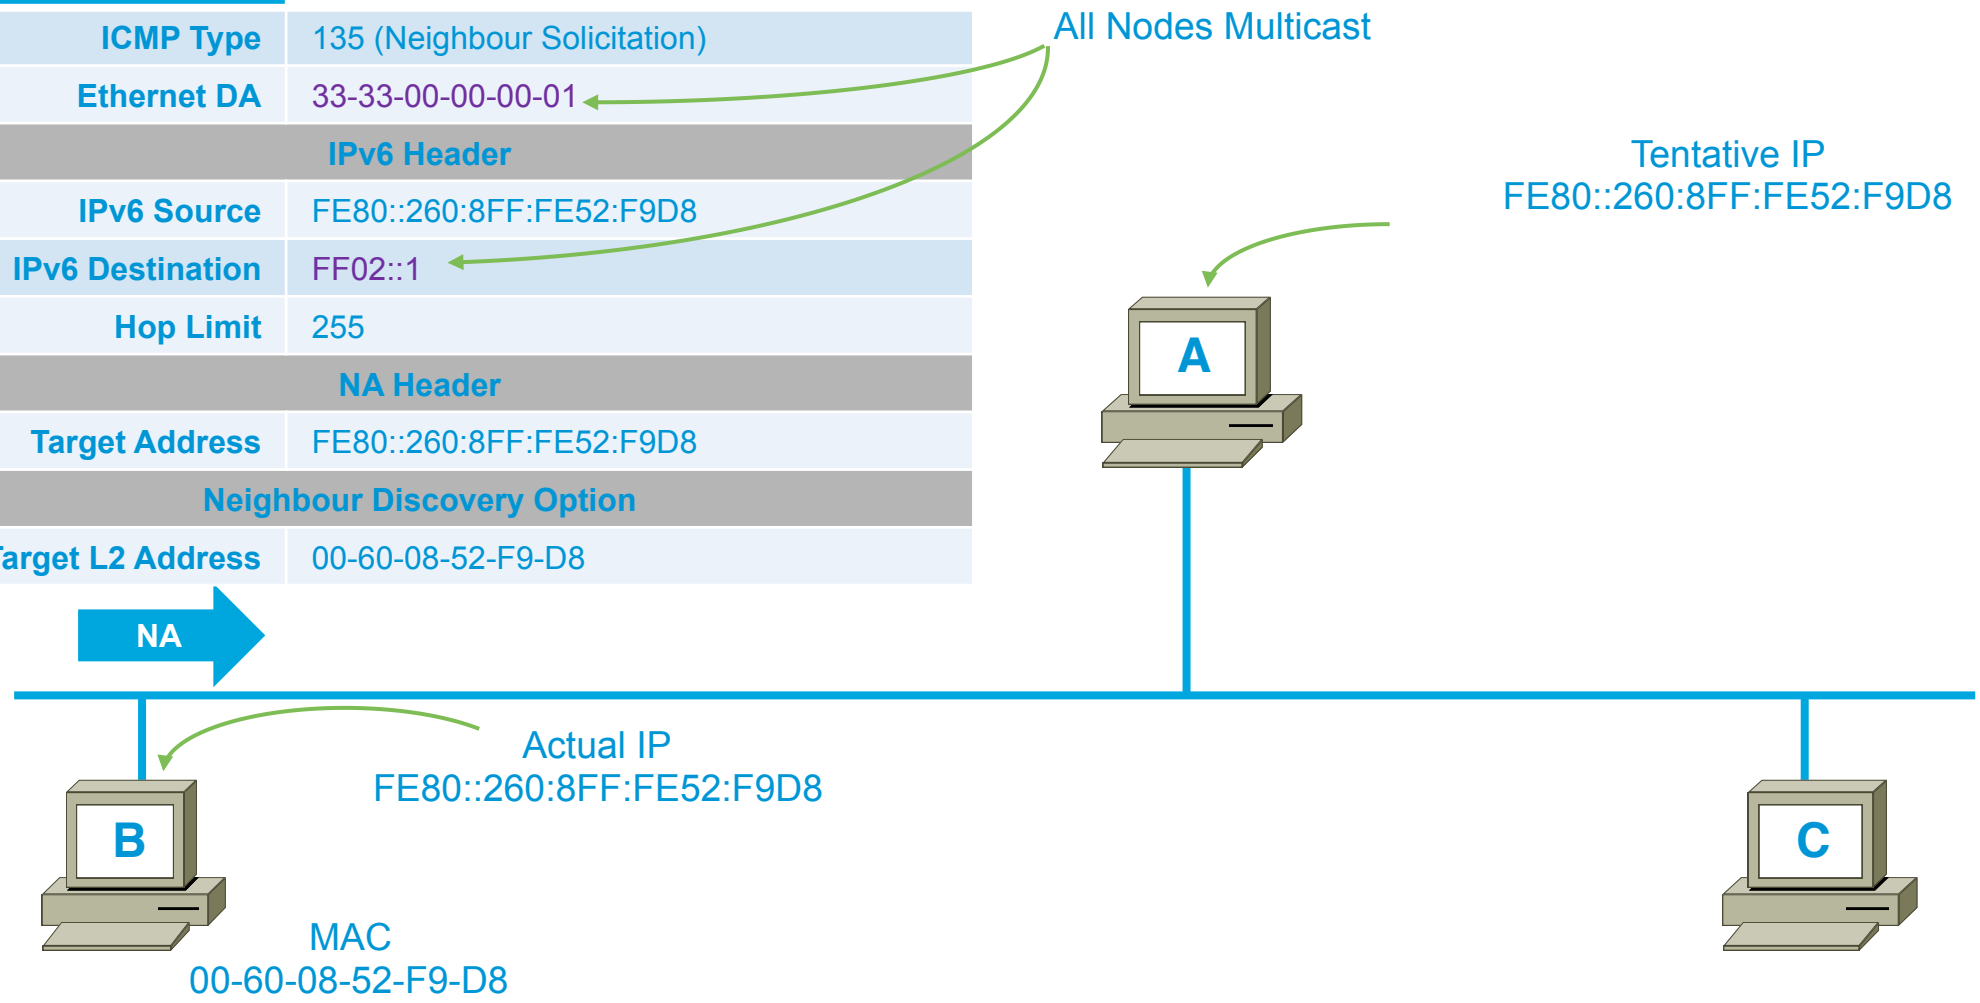

© 2011 Cisco and/or its affiliates. All rights reserved.

ICMPv6 Redirection

- Redirect is used by a router to informs hosts of a better first hop

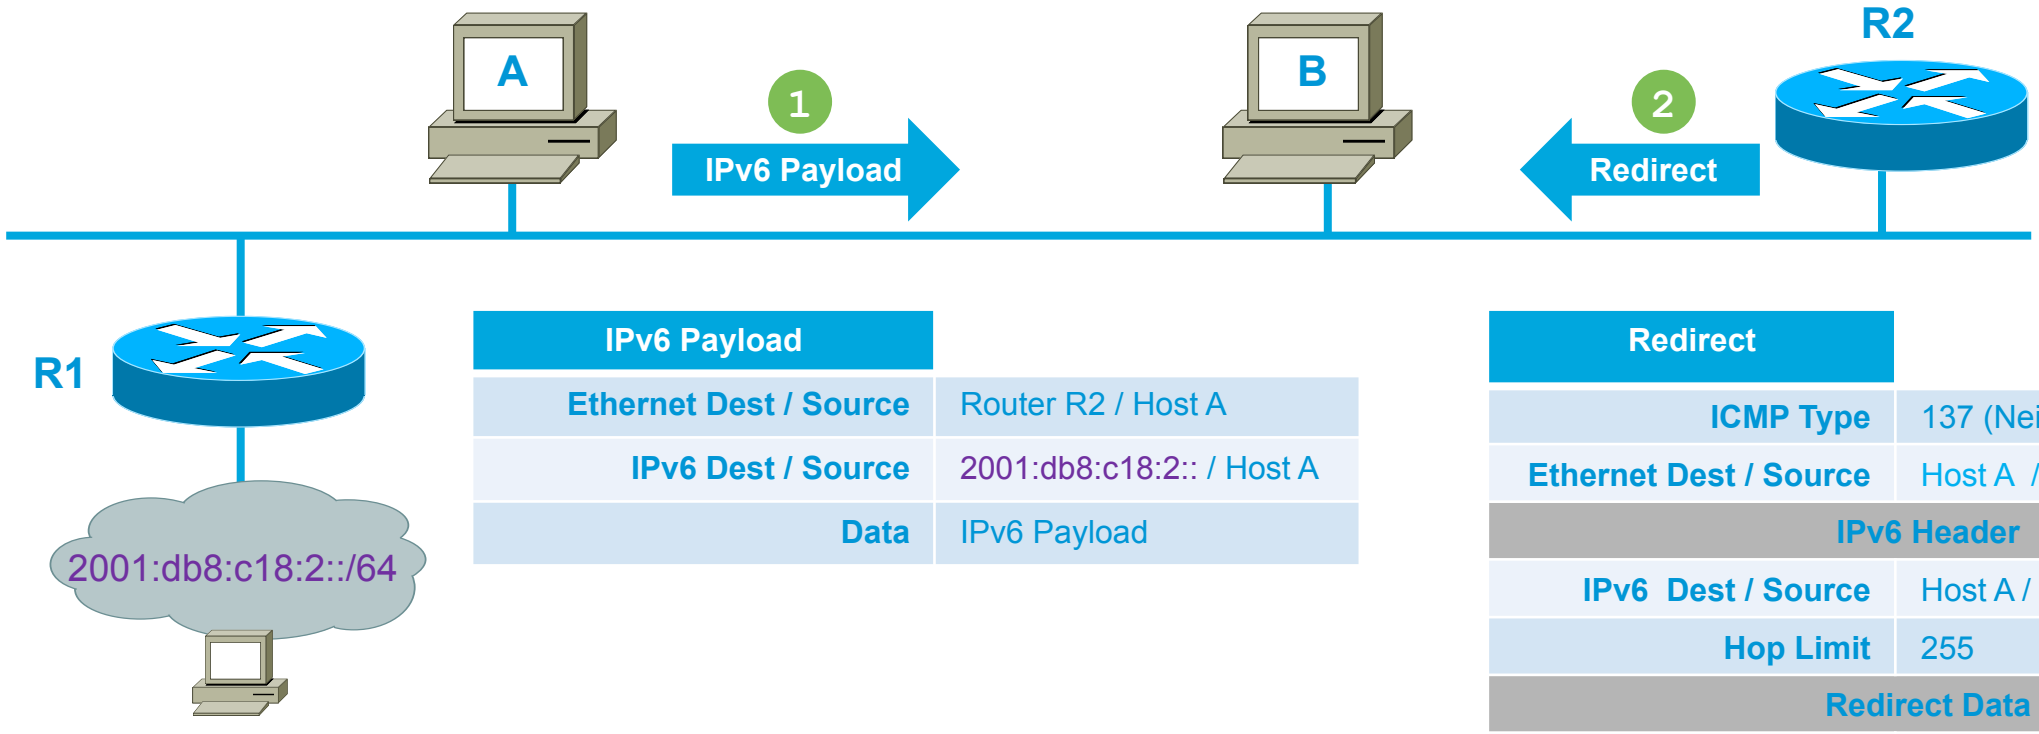

IPv6 Payload	
Ethernet Dest / Source	Router R2 / Host A
IPv6 Dest / Source	2001:db8:c18:2:: / Host A
Data	IPv6 Payload

Redirect	
ICMP Type	137 (Neighbour Solicitation)
Ethernet Dest / Source	Host A / Router R2
IPv6 Header	
IPv6 Dest / Source	Host A / Router R2
Hop Limit	255
Redirect Data	
Target Address	Router R1
Redirected Prefix	2001:db8:c18:2::/64

© 2011 Cisco and/or its affiliates. All rights reserved.

Stateless Address Autoconfiguration (RFC4862)

- Autoconfiguration is used to automatically assigned an address to a host “plug and play”
  - Generating a link-local address,
  - Generating global addresses via stateless address autoconfiguration
  - Duplicate Address Detection procedure to verify the uniqueness of the addresses on a link

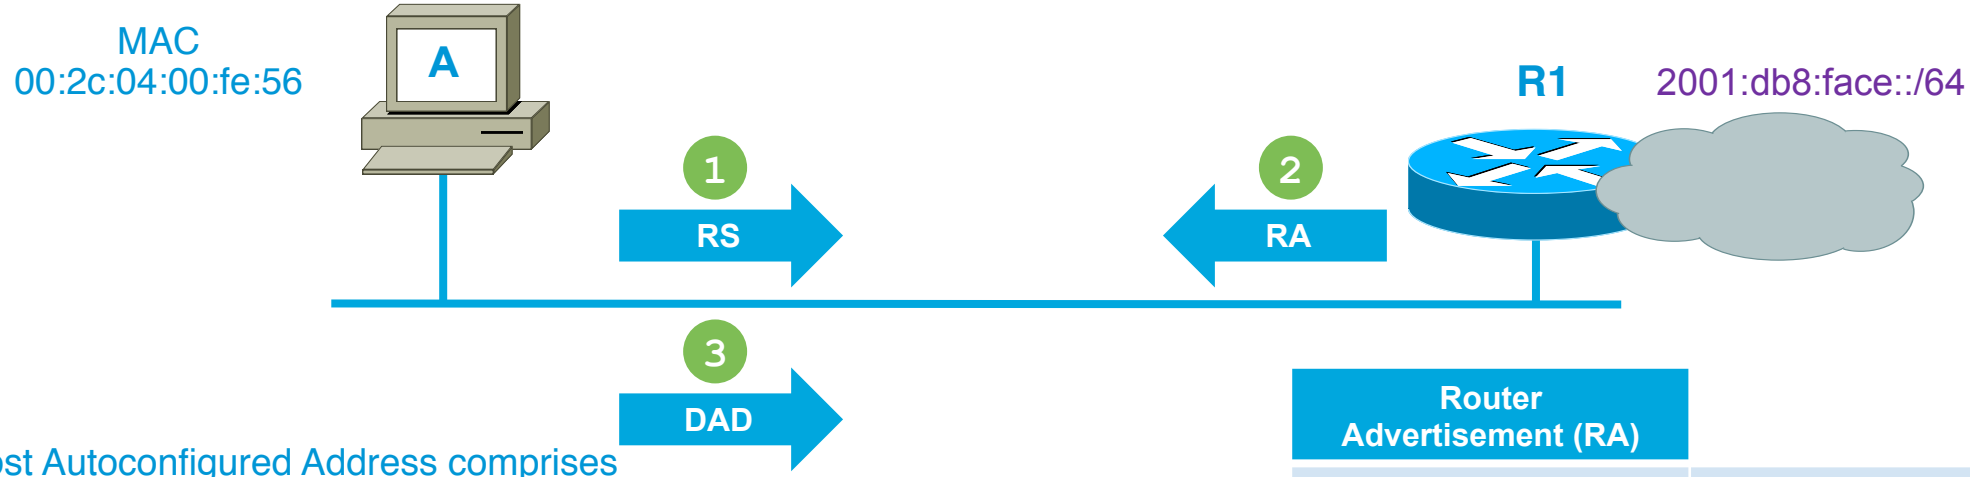

Host Autoconfigured Address comprises Prefix Received + Link-Layer Address if DAD check passes

```
2001:db8:face::22c:4ff:fe00:fe56
```

Router Advertisement (RA)	
Ethernet DA/SA	Router R2 / Host A
Prefix Information	2001:db8:face::/64
Default Router	Router R1

© 2011 Cisco and/or its affiliates. All rights reserved.

Prefix Renumbering

- Prefixes can be given a lifetime in RA messages
- Allows seamless transition for renumbering to a new prefix

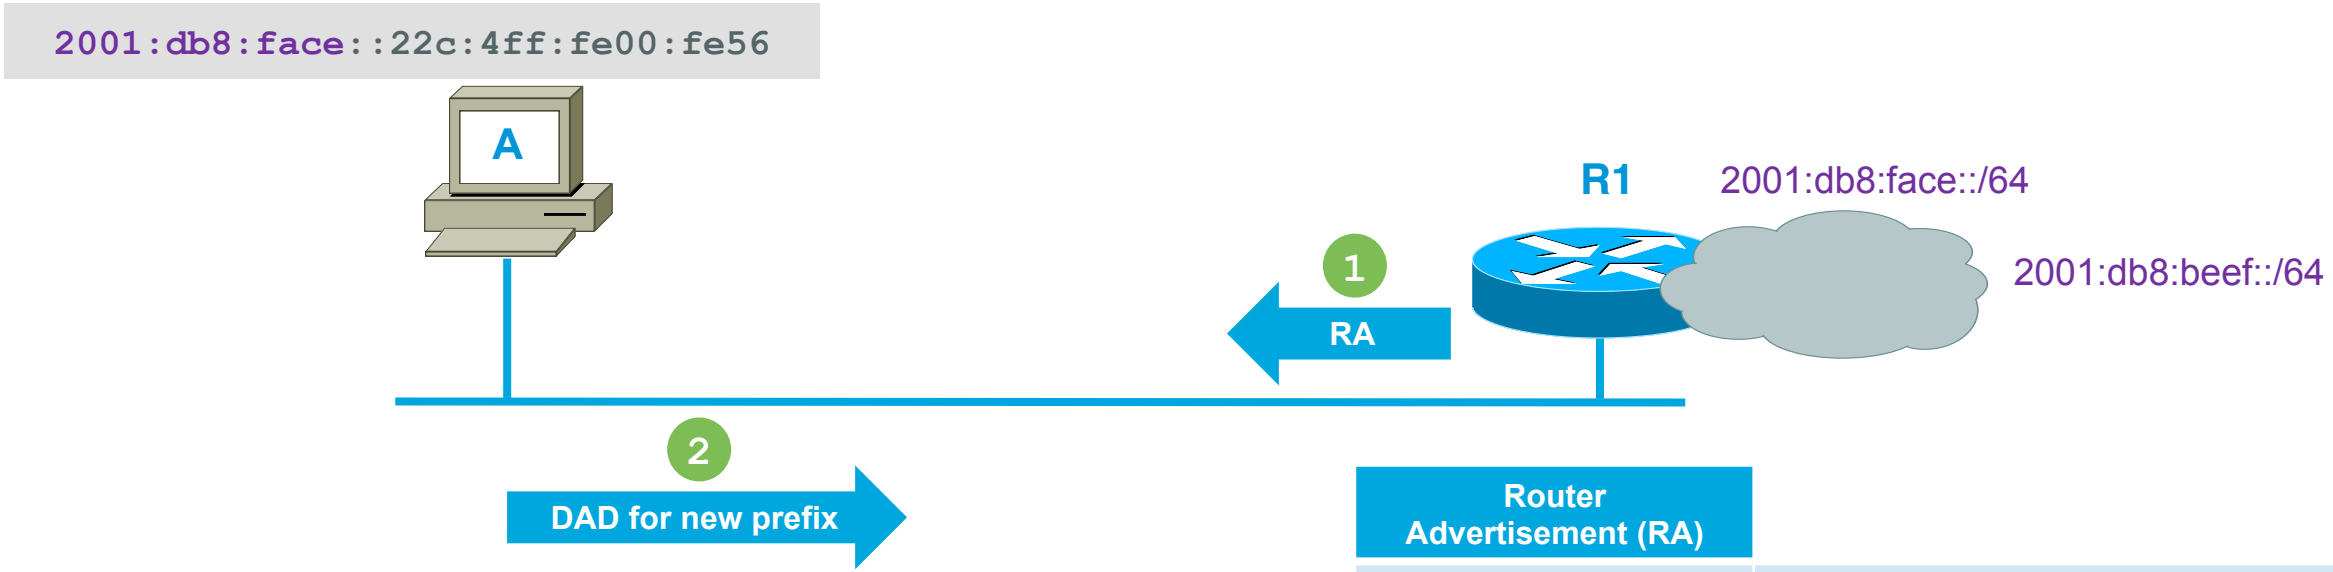

New Prefix

2001:db8:beef::22c:4ff:fe00:fe56

Router Advertisement (RA)	
Ethernet DA/SA	Router R2 / Host A
Current Prefix	2001:db8:face::/64, Lifetime 30 seconds
New Prefix	2001:db8:beef::/64, Lifetime 30 seconds

© 2011 Cisco and/or its affiliates. All rights reserved.

Default Router Selection

```
!
interface Ethernet0/0
 ipv6 nd reachable-time 15000
 ipv6 nd router-preference Low
!
```

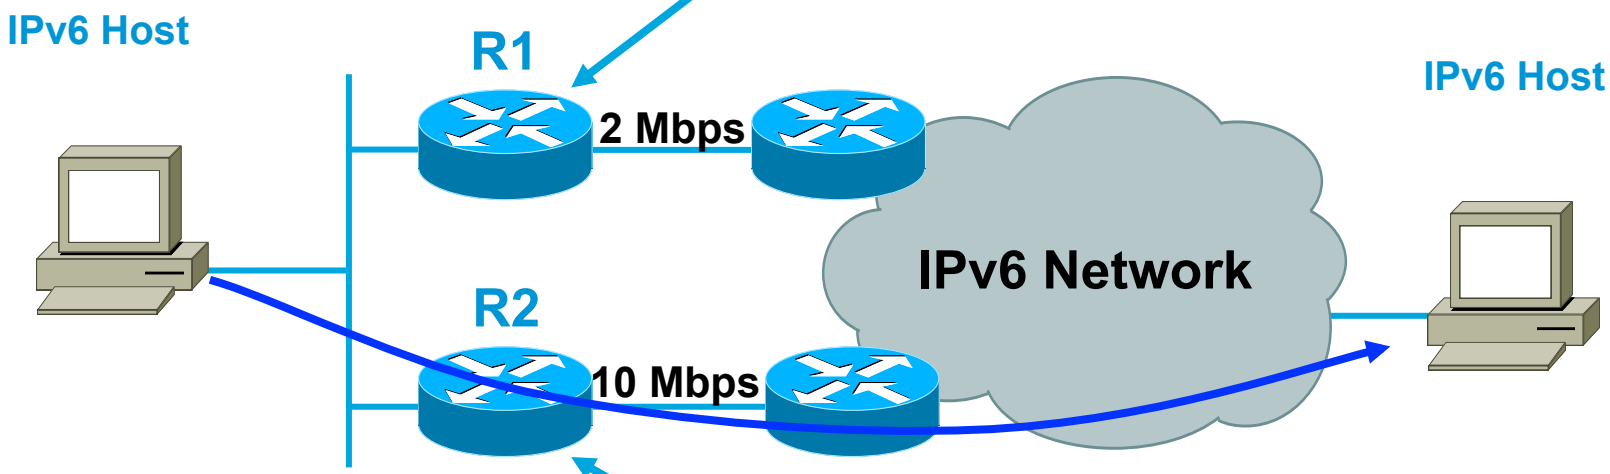


```
!
interface Ethernet0/0
 ipv6 nd reachable-time 15000
 ipv6 nd router-preference High
!
```

© 2011 Cisco and/or its affiliates. All rights reserved.

Default Router Selection

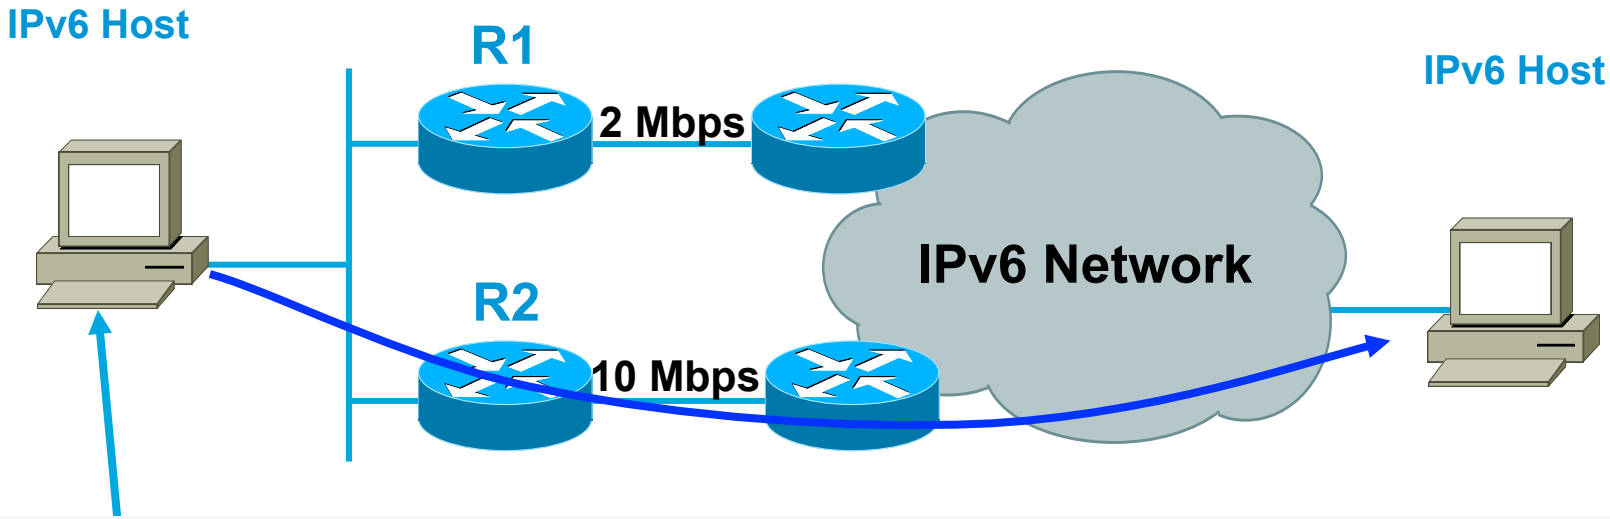


```
R200#sh ipv6 router
Router FE80::A8BB:CCFF:FE00:CA00 on Ethernet0/0, last update 0 min
  Hops 64, Lifetime 1800 sec, AddrFlag=0, OtherFlag=0, MTU=1500
  HomeAgentFlag=0, Preference=High
  Reachable time 15000 msec, Retransmit time 0 msec
  Prefix 2001:1::/64 onlink autoconfig
    Valid lifetime 2592000, preferred lifetime 604800
Router FE80::A8BB:CCFF:FE00:C900 on Ethernet0/0, last update 2 min
  Hops 64, Lifetime 1800 sec, AddrFlag=0, OtherFlag=0, MTU=1500
  HomeAgentFlag=0, Preference=Low
  Reachable time 15000 msec, Retransmit time 0 msec
  Prefix 2001:1::/64 onlink autoconfig
    Valid lifetime 2592000, preferred lifetime 604800
```

© 2011 Cisco and/or its affiliates. All rights reserved.

CISCO

Introduction to IPv6 Routing

Cisco live!

Agenda

- Introduction
- OSPFv3
- IS-IS for IPv6
- BGP for IPv6
- EIGRP for IPv6
- Co-existence
- Other Considerations
- Conclusions

© 2014 Cisco and/or its affiliates. All rights reserved.

Introduction

Questions You Might Have?

- Do I need an IPv6 routing protocol?
  - Dual-stack vs. other transition mechanisms
- Which IPv6 routing protocol should I choose?
- Do I have to start from scratch to learn a new IPv6 routing protocol?
- What are the main similarities and differences between the IPv6 routing protocols and their IPv4 counterparts?
- What practical design and deployment points should I consider?
- What about co-existence of IPv4 and IPv6 routing protocols?

© 2014 Cisco and/or its affiliates. All rights reserved.

IPv6 and IPv4 Routing Protocols

RIP	RIPv2 for IPv4 RIPng for IPv6 Distinct but similar protocols with RIPng taking advantage of IPv6 specificities
OSPF	OSPFv2 for IPv4 OSPFv3 for IPv6 Distinct but similar protocols with OSPFv3 being a cleaner implementation that takes advantage of IPv6 specificities
IS-IS	Extended to support IPv6 Natural fit to some of the IPv6 foundational concepts Supports Single and Multi Topology operation
EIGRP	Extended to support IPv6 Some changes reflecting IPv6 characteristics
BGP	Extended to support IPv6 through multi-protocol extensions

© 2014 Cisco and/or its affiliates. All rights reserved.

Dual-stack

- 2 ways of looking at dual-stack
  - From an end-point perspective:
    - IPv4 and IPv6 protocol stacks are both enabled.
    - The choice of IP version is determined by DNS lookup and application preference
  - From a network perspective:
    - IPv4 and IPv6 protocol stacks are both enabled
    - IPv4 and IPv6 IGP / BGP routing protocols are both enabled
    - The IPv4 and IPv6 control planes (routing protocol(s)) operate *mostly* independently
    - IPv4 and IPv6 packet forwarding operates independently
    - This is called “ships-in-the-night” behaviour

© 2014 Cisco and/or its affiliates. All rights reserved.

OSPFv3

OSPFv3

Implementation	Similar Concepts as OSPFv2: - Runs directly over IPv6 (Next Header 89) - Uses the same basic packet types - Neighbor discovery and adjacency formation mechanisms are identical (All OSPF Routers FF02::5, All OSPF DRs FF02::6) - LSA flooding and aging mechanisms are identical - Same interface types (P2P, P2MP, Broadcast, NBMA, Virtual) Independent process from OSPFv2
Important Differences	**OSPFv3 Is Running per Link Instead of per Node (and IP Subnet)** **Support of Multiple Instances per Link:** - New field (instance) in OSPF packet header allows running multiple instances per link - Instance ID should match before packet is being accepted - Useful for traffic separation, multiple areas per link **Generalization of Flooding Scope:** - Three flooding scopes for LSAs (link-local scope, area scope, AS scope) and they are coded in the LS type explicitly

© 2014 Cisco and/or its affiliates. All rights reserved.

OSPFv3

Important Differences (cont.)

Address Semantic Changes in LSA:
- Router and Network LSA carry only topology information
- Router LSA can be split across multiple LSAs; Link State ID in LSA header is a fragment ID
- Intra area prefixes are carried in a new LSA payload called intra-area-prefix-LSAs
- Prefixes are carried in the payload of inter-area and external LSA

Explicit Handling of Unknown LSA:
- The handling of unknown LSA is coded via U-bit in LS type
- When U bit is set, the LSA is flooded within the corresponding flooding scope, as if it was understood
- When U bit is not set, the LSA is flooded within the link local scope

Authentication Is Removed from OSPF:
- Authentication in OSPFv3 has been removed and OSPFv3 relies now on IPv6 authentication header since OSPFv3 runs over IPv6
- Autype and Authentication field in the OSPF packet header therefore have been suppressed

© 2014 Cisco and/or its affiliates. All rights reserved.

OSPFv3

Important Differences (cont.)	**OSPF Packet Format has Been Changed:** - The mask field has been removed from Hello packet - IPv6 prefixes are only present in payload of Link State update packet **Two New LSAs Have Been Introduced:** - Link-LSA has a link local flooding scope and has three purposes Carry IPv6 link local address used for NH calculation Advertise IPv6 global address to other routers on the link (used for multi-access link) Convey router options to DR on the link - Intra-area-prefix-LSA to advertise router's IPv6 address within the area
Notes	**Standardization**: Main standard: http://tools.ietf.org/html/rfc5340 Support for multiple AFs (multiple instance): http://tools.ietf.org/html/rfc5838

© 2014 Cisco and/or its affiliates. All rights reserved.

OSPFv3 Header

- Size of the header is reduced from 24 bytes to 16
- Router ID is still a 32 bit number uniquely identifying a router in the domain
- Instance ID is a new field that is used to have multiple OSPF process instances per link. In order for 2 instances to talk to each other they need to have the same Instance ID. By default it is 0 and for any additional instance it is increased, Instance ID has local link significance only
- Authentication fields have been suppressed

OSPFv2

Version	Type	Packet Length
Router ID		
Area ID		
Checksum		Autype
Authentication		
Authentication		

OSPFv3

Version	Type	Packet Length
Router ID		
Area ID		
Checksum	Instance ID	0

© 2014 Cisco and/or its affiliates. All rights reserved.
Cisco Public

OSPFv3 LSA Header

- All LSAs begin with a common 20 byte header (like OSPFv2) – LS type field expanded
- LS Age: The time in seconds since the LSA was originated
- Link State ID: This field identifies the piece of the routing domain that is being described by the LSA. Depending on the LSA's LS type, the Link State ID takes on its value.
- The behavior of assigning this value has changed from OSPFv2 to OSPFv3
- Advertising Router: ID of the router originating the packet.

OSPFv2

LS age	Options	LS type
Link State ID		
Advertising Router		
LS sequence numer		
LS checksum	Length	

OSPFv3

LS age	LS type
Link State ID	
Advertising Router	
LS sequence numer	
LS checksum	Length

© 2014 Cisco and/or its affiliates. All rights reserved.
Cisco Public

OSPFv3 LSA Types

	LSA Function Code	LSA Type Code
Router-LSA	1	0x2001
Network-LSA	2	0x2002
Inter-Area-Prefix-LSA	3	0x2003
Inter-Area-Router-LSA	4	0x2004
AS-External-LSA	6	0x4005
NSSA-LSA	7	0x2007
Link-LSA	8	0x0008
Intra-Area-Prefix-LSA	9	0x2009

- LSA Type Code encodes “Flooding Scope”

S1	S2	Flooding Scope
0	0	Link-Local
0	1	Area
1	0	AS
1	1	Reserved

© 2014 Cisco and/or its affiliates. All rights reserved.

OSPFv3 vs. OSPFv2 Summary

	OSPFv2 (RFC 2328)	OSPFv3 (RFC 5340)
Type	Link state	Link state
Metric	Cost (1-65535)	Cost (1-65535)
Loop prevention	Dijkstra	Dijkstra
Administrative distance	110	110
L3 transport S/D	IPv4 unicast / 224.0.0.5 IPv4 unicast / 224.0.0.6	IPv6 LL unicast / FF02::5 IPv6 LL unicast / FF02::6
L4 transport	IP 89	Next Header 89

```
interface Loopback0
 ip address 10.10.0.1 255.255.255.255
!
 interface Ethernet0/0
 ip address 10.0.0.1 255.255.255.0
!
router ospf 1
 network 10.0.0.0 0.255.255.255 area 0
 passive-interface loopback0
```

```
ipv6 unicast-routing
!
interface Loopback0
 ipv6 address 2001:DB8:1000::1/128
 ipv6 ospf 1 area 0
!
interface Ethernet0/0
 ipv6 address 2001:DB8::1/64
 ipv6 ospf 1 area 0
!
ipv6 router ospf 1
 router-id 10.10.10.1
 passive-interface loopback0
```

For Your Reference

© 2014 Cisco and/or its affiliates. All rights reserved.

OSPFv2 / OSPFv3 Sample Topology

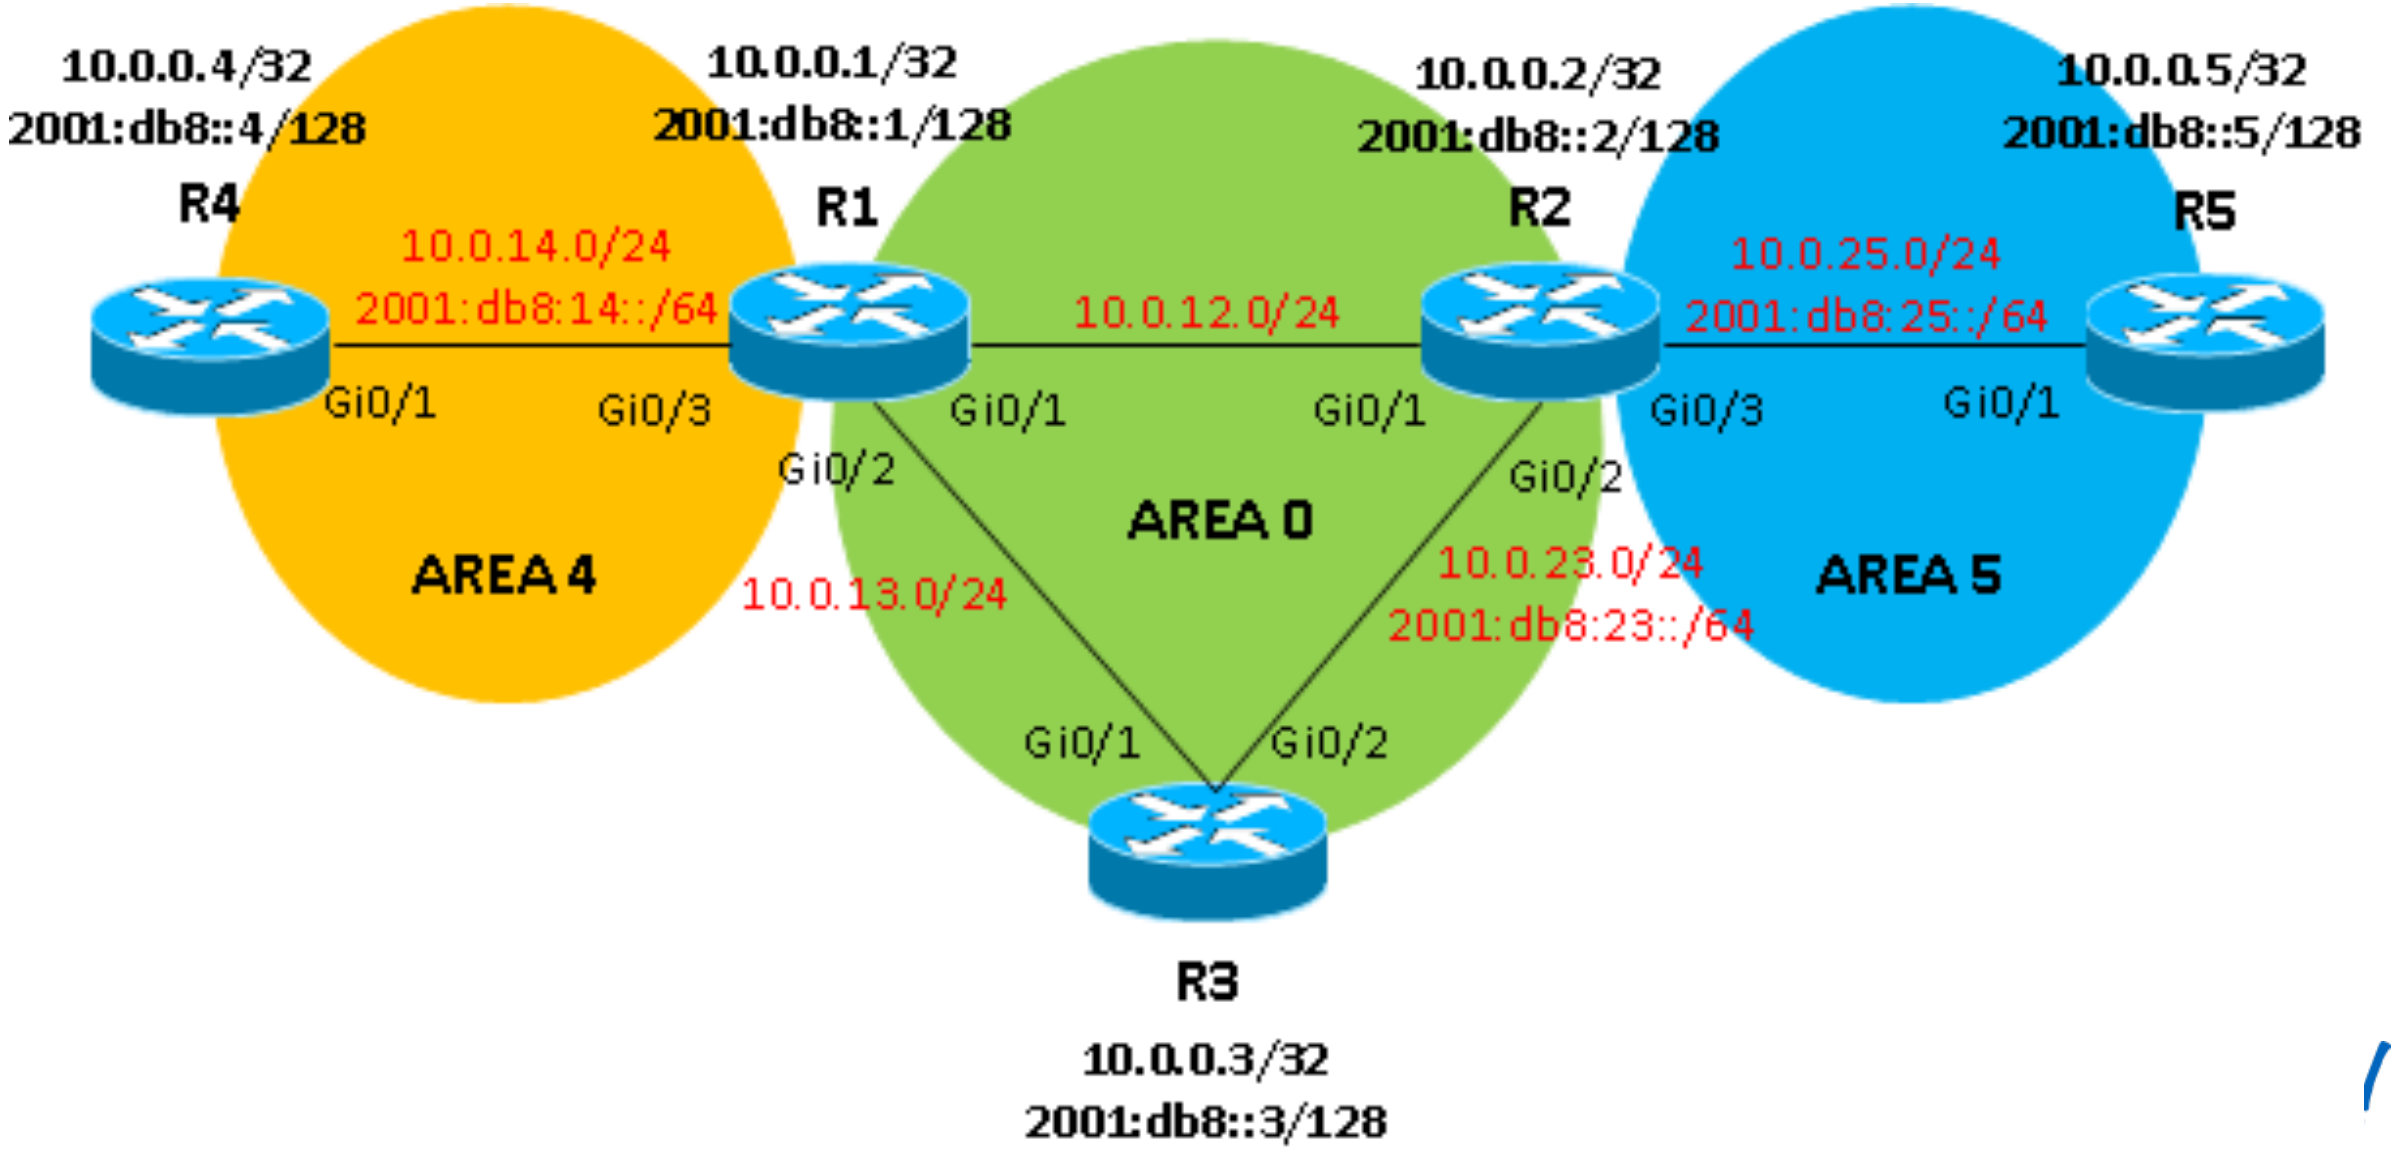

© 2014 Cisco and/or its affiliates. All rights reserved.

OSPFv3 Configuration / Interfaces

```
!
ipv6 unicast-routing
!
interface Loopback0
```

```
R1#show ipv6 ospf interface gigabitEthernet 0/1
GigabitEthernet0/1 is up, line protocol is up
  Link Local Address FE80::F816:3EFF:FE59:54FE, Interface ID 3
  Area 0, Process ID 1, Instance ID 0, Router ID 10.0.0.1
  Network Type POINT_TO_POINT, Cost: 1
  Transmit Delay is 1 sec, State POINT TO POINT
  Timer intervals
    Hello due in
  Graceful restar
  Index 1/2/2, f
  Next 0x0(0)/0x(
  Last flood scar
  Last flood scar
  Neighbor Count
    Adjacent wit
  Suppress hello
R1#
```

```
R3#sh ipv6 ospf interface gigabitEthernet 0/2
GigabitEthernet0/2 is up, line protocol is up
  Link Local Address FE80::F816:3EFF:FEF1:DD7B, Interface ID 4
  Area 0, Process ID 1, Instance ID 0, Router ID 10.0.0.3
  Network Type BROADCAST, Cost: 1
  Transmit Delay is 1 sec, State BDR, Priority 1
  Designated Router (ID) 10.0.0.2, local address FE80::F816:3EFF:FE86:CC18
  Backup Designated router (ID) 10.0.0.3, local address
FE80::F816:3EFF:FEF1:DD7B
  Timer intervals configured, Hello 10, Dead 40, Wait 40, Retransmit 5
    Hello due in 00:00:01
  Graceful restart helper support enabled
  Index 1/3/3, flood queue length 0
  Next 0x0(0)/0x0(0)/0x0(0)
  Last flood scan length is 1, maximum is 5
  Last flood scan time is 0 msec, maximum is 0 msec
  Neighbor Count is 1, Adjacent neighbor count is 1
    Adjacent with neighbor 10.0.0.2  (Designated Router)
  Suppress hello for 0 neighbor(s)
R3#
```

```
 ne    ND DAD is en
 ne    ND reachable
 ne    ND advertise
!      ND advertise
ipv    ND router ad
 rou   ND router ad
!      ND advertise
       Hosts use st
R1#
```

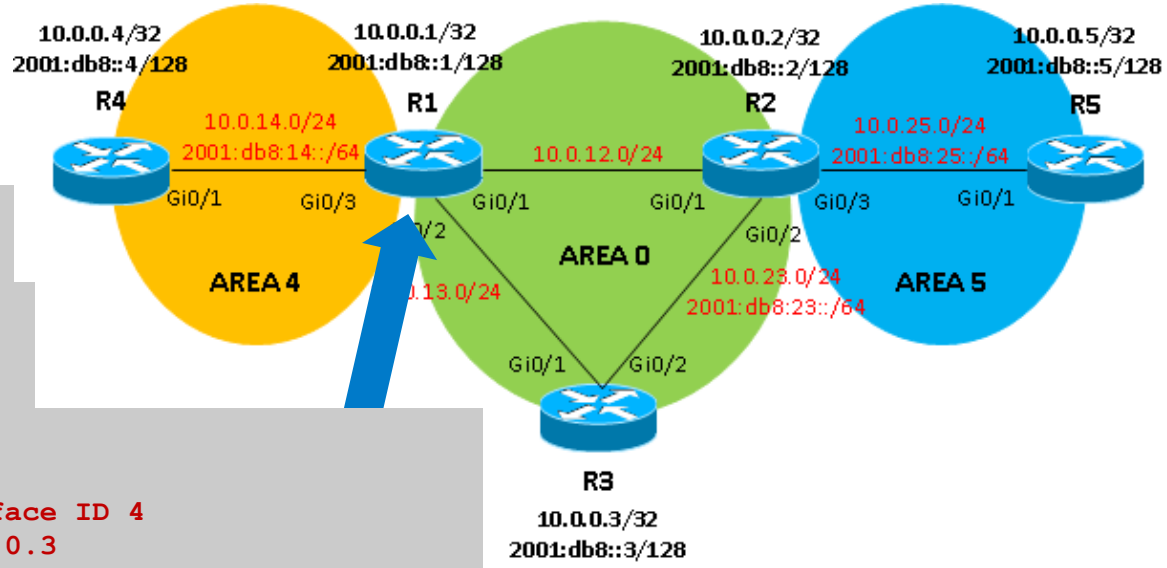

© 2014 Cisco and/or its affiliates. All rights reserved.

OSPFv3 Neighbours

```
R1#show ipv6 ospf neighbor gigabitEthernet 0/1 detail

            OSPFv3 Router with ID (10.0.0.1) (Process ID 1)

 Neighbor 10.0.0.2
    In the area 0 via interface GigabitEthernet0/1
    Neighbor: interface-id 3, link-local address
FE80::F816:3EFF:FE82:591
    Neighbor priority is 0, State is FULL, 6 state changes
    Options is 0x000013 in Hello (V6-Bit, E-Bit, R-bit)
    Options is 0x000013 in DBD (V6-Bit, E-Bit, R-bit)
    Dead timer due in 00:00:32
    Neighbor is up for 00:09:45
    Index 1/1/1, retransmission queue length 0, number of retransmission
0
    Fir
    Las
    Las
R1#
```

```
R3#show ipv6 ospf neighbor gigabitEthernet 0/2 detail

            OSPFv3 Router with ID (10.0.0.3) (Process ID 1)

 Neighbor 10.0.0.2
    In the area 0 via interface GigabitEthernet0/2
    Neighbor: interface-id 4, link-local address FE80::F816:3EFF:FE6D:10E1
    Neighbor priority is 1, State is FULL, 6 state changes
    DR is 10.0.0.2 BDR is 10.0.0.3
    Options is 0x000013 in Hello (V6-Bit, E-Bit, R-bit)
    Options is 0x000013 in DBD (V6-Bit, E-Bit, R-bit)
    Dead timer due in 00:00:38
    Neighbor is up for 00:11:13
    Index 1/2/2, retransmission queue length 0, number of retransmission 0
    First 0x0(0)/0x0(0)/0x0(0) Next 0x0(0)/0x0(0)/0x0(0)
    Last retransmission scan length is 0, maximum is 0
    Last retransmission scan time is 0 msec, maximum is 0 msec
R3#
```

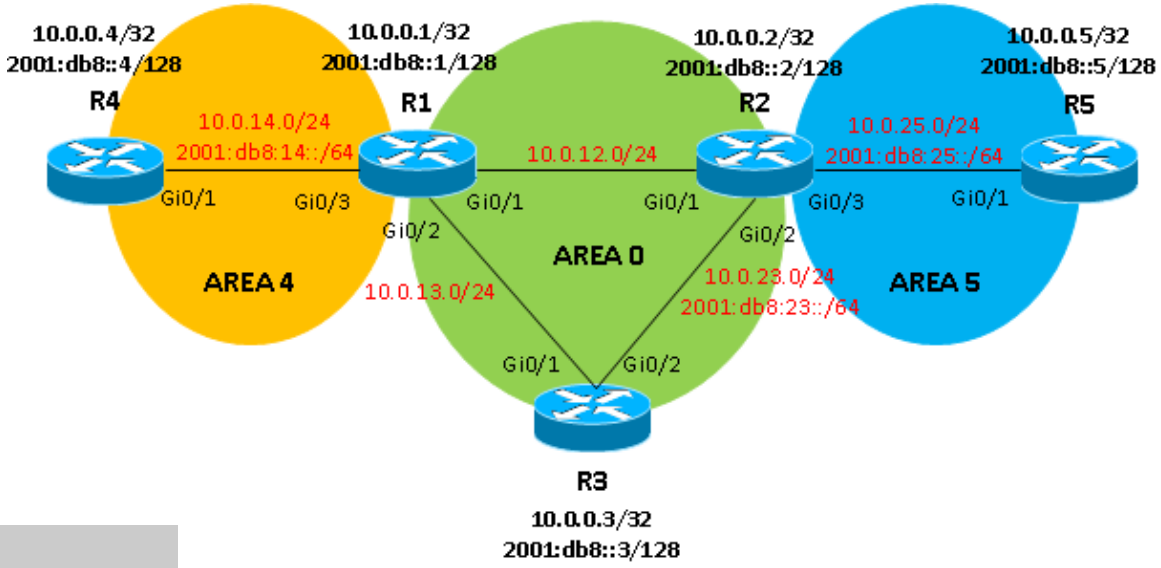

OSPFv3 RIB

```
R1#show ipv6 route
IPv6 Routing Table - default - 6 entries
Codes: C - Connected, L - Local, S - Static, U - Per-user Static
route
       B - BGP, HA - Home Agent, MR - Mobile Router, R - RIP
       H - NHRP, I1 - ISIS L1, I2 - ISIS L2, IA - ISIS interarea
       IS - ISIS summary, D - EIGRP, EX - EIGRP external, NM -
NEMO
       ND - ND Default, NDp - ND Prefix, DCE - Destination, NDr -
Redirect
       O - OSPF Intra, OI - OSPF Inter, OE1 - OSPF ext 1, OE2 -
OSPF ext 2
       ON1 - OSPF NSSA ext 1, ON2 - OSPF NSSA ext 2, ls - LISP
site
       ld - LISP dyn-EID, a - Application
LC  2001:DB8::1/128 [0/0]
     via Loopback0, receive
O   2001:DB8::2/128 [110/1]
     via FE80::F816:3EFF:FE82:591, GigabitEthernet0/1
O   2001:DB8::3/128 [110/1]
     via FE80::F816:3EFF:FED7:C7D7, GigabitEthernet0/2
O   2001:DB8::4/128 [110/1]
     via FE80::F816:3EFF:FE18:9E63, GigabitEthernet0/3
OI  2001:DB8::5/128 [110/2]
     via FE80::F816:3EFF:FE82:591, GigabitEthernet0/1
L   FF00::/8 [0/0]
     via Null0, receive
R1#
```

```
R1#show
* 2001:
    tag
* 2001:
    tag
    via
* 2001:
    tag
    via
* 2001:
    tag
    via
* 2001:
    tag
    via
    LSA
R1#
```

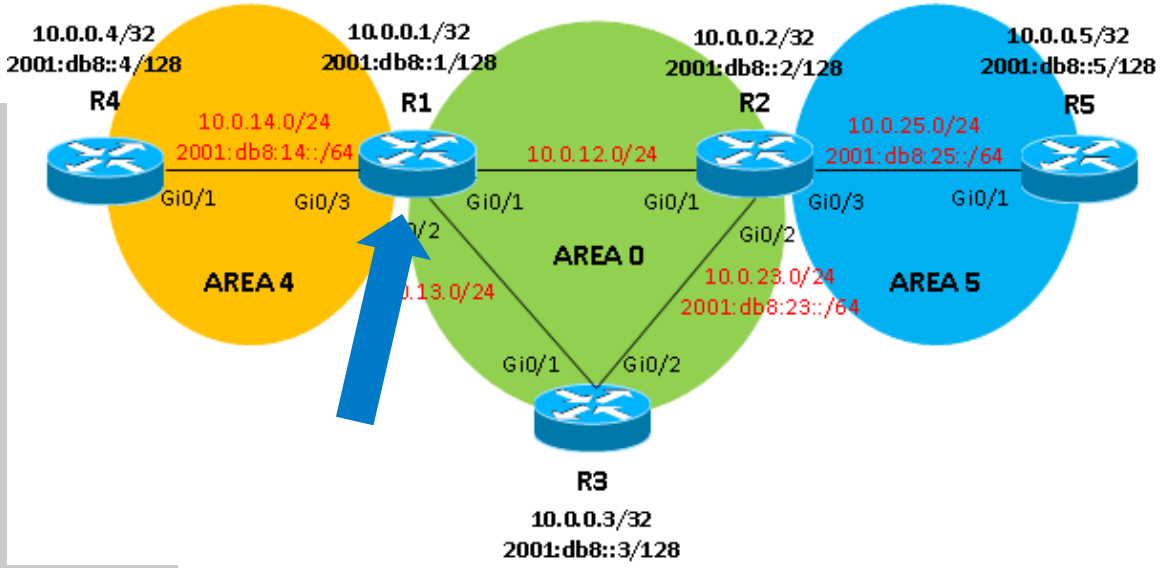

© 2014 Cisco and/or its affiliates. All rights reserved.

OSPFv3 LSDB

```
R1#show ipv6 ospf database database-summary

            OSPFv3 Router with ID (10.0.0.1) (Process ID 1)

Area 0 database summary
  LSA Type              Count     Delete    Maxage
  Router                3         0         0
  Network               1         0         0
  Link                  4         0         0
  Prefix                3         0         0
  Inter-area Prefix     2         0         0
  Inter-area Router     0         0         0
  Type-7 External       0         0         0
  Unknown               0         0         0
  Subtotal              13        0         0

Area 4 database summary
  LSA Type              Count     Delete    Maxage
  Router                2         0         0
  Network               0         0         0
  Link                  2         0         0
  Prefix                1         0         0
  Inter-area Prefix     4         0         0
  Inter-area Router     0         0         0
  Type-7 External       0         0         0
  Unknown               0         0         0
  Subtotal              9         0         0
...
R1#
```

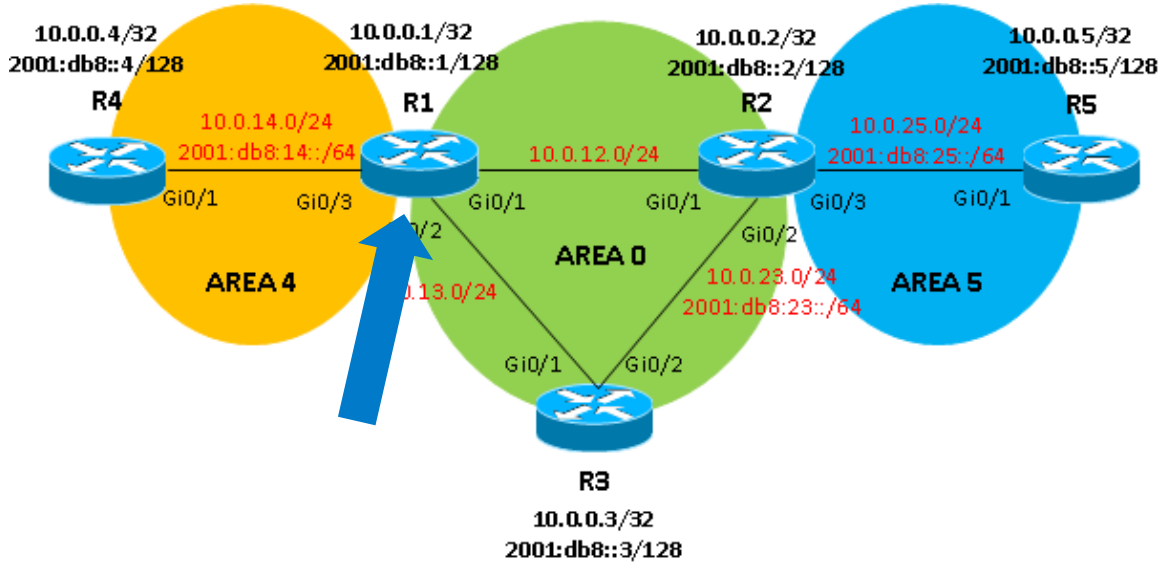

© 2014 Cisco and/or its affiliates. All rights reserved.

OSPFv3 Router LSA

```
R1#show ipv6 ospf database router adv-router 10.0.0.2

            OSPFv3 Router with ID (10.0.0.1) (Process ID 1)

                Router Link States (Area 0)

  Routing Bit Set on this LSA
  LS age: 957
  Options: (V6-Bit, E-Bit, R-bit, DC-Bit)
  LS Type: Router Links
  Link State ID: 0
  Advertising Router: 10.0.0.2
  LS Seq Number: 8000000D
  Checksum: 0xAEFF
  Length: 56
  Area Border Router
  Number of Links: 2

    Link connected to: a Transit Network
      Link Metric: 1
      Local Interface ID: 4
      Neighbor (DR) Interface ID: 4
      Neighbor (DR) Router ID: 10.0.0.2

    Link connected to: another Router (point-to-point)
      Link Metric: 1
      Local Interface ID: 3
      Neighbor Interface ID: 3
      Neighbor Router ID: 10.0.0.1


R1#
```

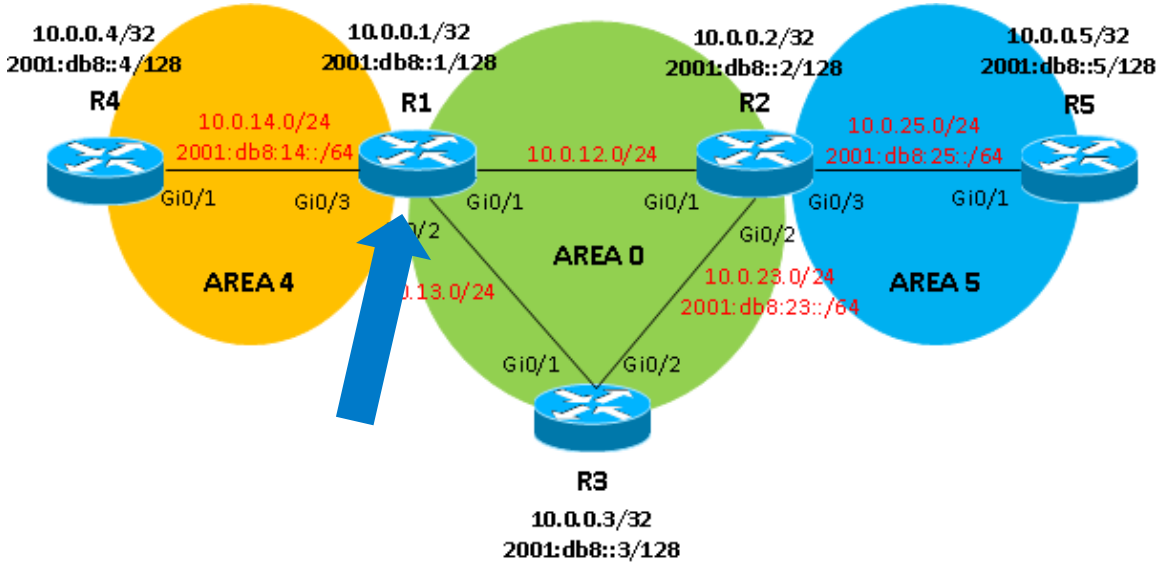

rights reserved.

OSPFv3 Network LSA

```
R1#show ipv6 ospf database network adv-router 10.0.0.2

            OSPFv3 Router with ID (10.0.0.1) (Process ID 1)

                Net Link States (Area 0)

  LS age: 1067
  Options: (V6-Bit, E-Bit, R-bit, DC-Bit)
  LS Type: Network Links
  Link State ID: 4 (Interface ID of Designated Router)
  Advertising Router: 10.0.0.2
  LS Seq Number: 8000000B
  Checksum: 0x4194
  Length: 32
        Attached Router: 10.0.0.2
        Attached Router: 10.0.0.3

R1#
```

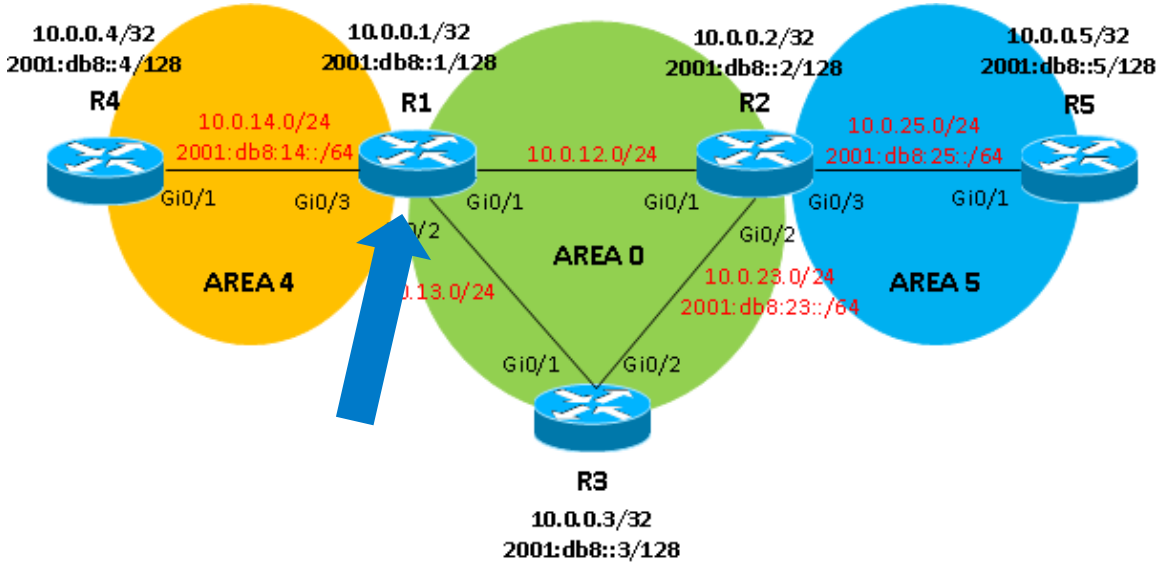

© 2014 Cisco and/or its affiliates. All rights reserved.

OSPFv3 Link LSA

```
R1#show ipv6 ospf database link adv-router 10.0.0.1

            OSPFv3 Router with ID (10.0.0.1) (Process ID 1)

                Link (Type-8) Link States (Area 0)

  LS age: 907
  Options: (V6-Bit, E-Bit, R-bit, DC-Bit)
  LS Type: Link-LSA (Interface: GigabitEthernet0/2)
  Link State ID: 4 (Interface ID)
  Advertising Router: 10.0.0.1
  LS Seq Number: 8000000C
  Checksum: 0xD637
  Length: 44
  Router Priority: 1
  Link Local Address: FE80::F816:3EFF:FEF4:2DFF
  Number of Prefixes: 0

  LS age: 907
  Options: (V6-Bit, E-Bit, R-bit, DC-Bit)
  LS Type: Link-LSA (Interface: GigabitEthernet0/1)
  Link State ID: 3 (Interface ID)
  Advertising Router: 10.0.0.1
  LS Seq Number: 8000000C
  Checksum: 0xBEE4
  Length: 44
  Router Priority: 1
  Link Local Address: FE80::F816:3EFF:FE48:E361
  Number of Prefixes: 0
....
```

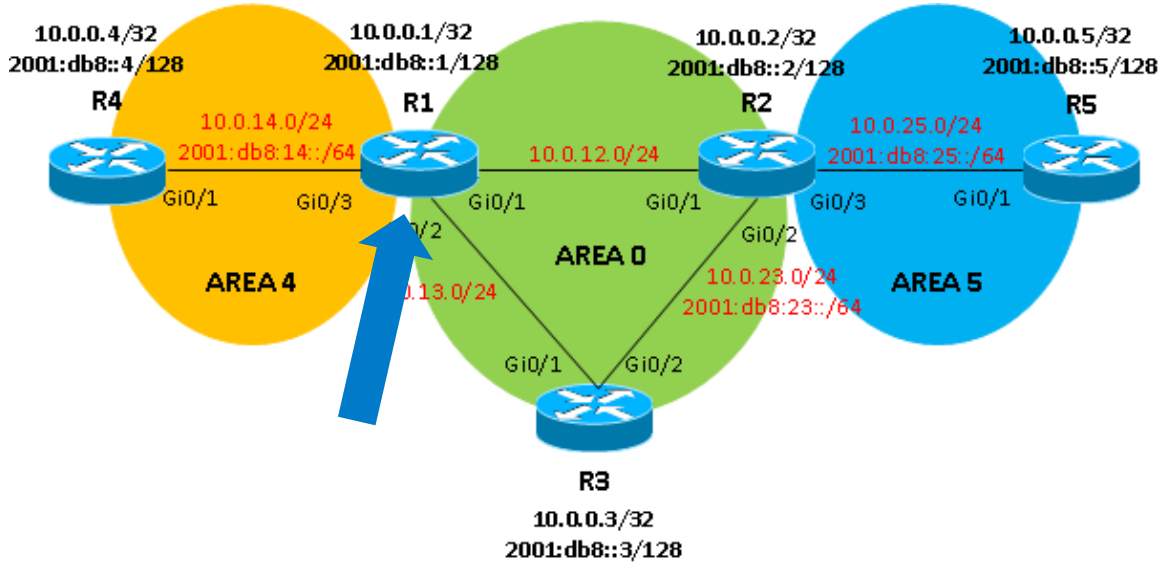


```
....
Link (Type-8) Link States (Area 4)

  LS age: 907
  Options: (V6-Bit, E-Bit, R-bit, DC-Bit)
  LS Type: Link-LSA (Interface: GigabitEthernet0/3)
  Link State ID: 5 (Interface ID)
  Advertising Router: 10.0.0.1
  LS Seq Number: 8000000C
  Checksum: 0xDE3A
  Length: 44
  Router Priority: 1
  Link Local Address: FE80::F816:3EFF:FE97:CCB1
  Number of Prefixes: 0

R1#
```


© 2014 Cisco and/or its affiliates. All rights reserved.

OSPFv3 Intra-Area Prefix LSA

```
R1#show ipv6 ospf database prefix adv-router 10.0.0.2

            OSPFv3 Router with ID (10.0.0.1) (Process ID 1)

                Intra Area Prefix Link States (Area 0)

  Routing Bit Set on this LSA
  LS age: 170
  LS Type: Intra-Area-Prefix-LSA
  Link State ID: 0
  Advertising Router: 10.0.0.2
  LS Seq Number: 8000000D
  Checksum: 0x8EE0
  Length: 52
  Referenced LSA Type: 2001
  Referenced Link State ID: 0
  Referenced Advertising Router: 10.0.0.2
  Number of Prefixes: 1
  Prefix Address: 2001:DB8::2
  Prefix Length: 128, Options: LA, Metric: 0
....
```

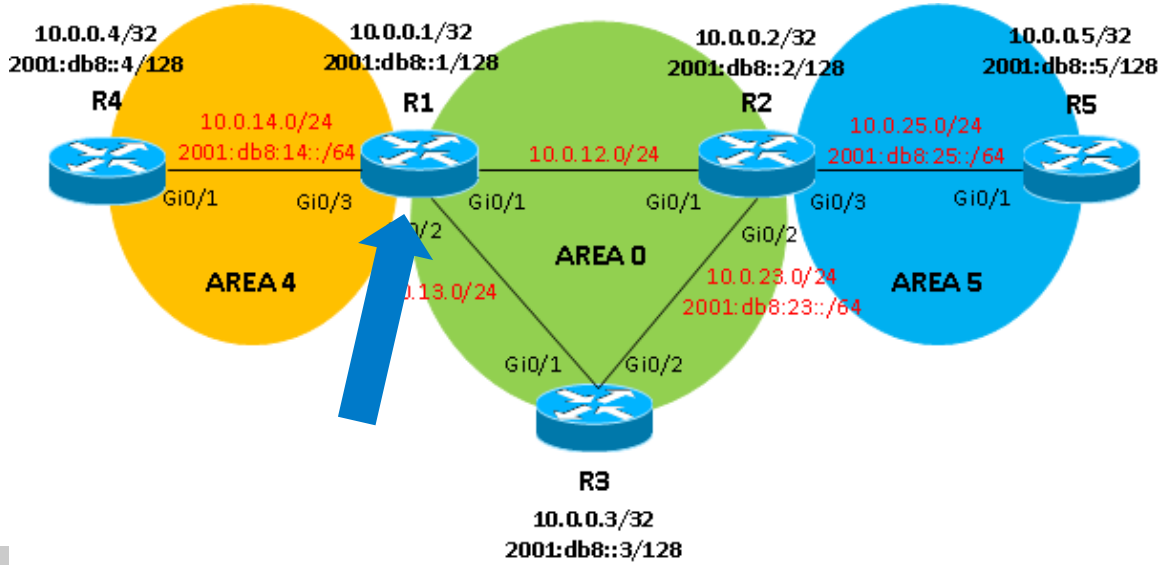


```
....
  Routing Bit Set on this LSA
  LS age: 75
  LS Type: Intra-Area-Prefix-LSA
  Link State ID: 4096
  Advertising Router: 10.0.0.2
  LS Seq Number: 80000001
  Checksum: 0x3758
  Length: 44
  Referenced LSA Type: 2002
  Referenced Link State ID: 4
  Referenced Advertising Router: 10.0.0.2
  Number of Prefixes: 1
  Prefix Address: 2001:DB8:23::
  Prefix Length: 64, Options: None, Metric: 0

R1#
```


© 2014 Cisco and/or its affiliates. All rights reserved.

OSPFv3 Inter-Area Prefix LSA

```
R1#show ipv6 ospf database inter-area prefix adv-router 10.0.0.2

            OSPFv3 Router with ID (10.0.0.1) (Process ID 1)

                Inter Area Prefix Link States (Area 0)

  Routing Bit Set on this LSA
  LS age: 115
  LS Type: Inter Area Prefix Links
  Link State ID: 0
  Advertising Router: 10.0.0.2
  LS Seq Number: 8000000D
  Checksum: 0x8623
  Length: 44
  Metric: 1
  Prefix Address: 2001:DB8::5
  Prefix Length: 128, Options: None

  Routing Bit Set on this LSA
  LS age: 1320
  LS Type: Inter Area Prefix Links
  Link State ID: 1
  Advertising Router: 10.0.0.2
  LS Seq Number: 80000001
  Checksum: 0x924A
  Length: 36
  Metric: 1
  Prefix Address: 2001:DB8:25::
  Prefix Length: 64, Options: None

R1#
```

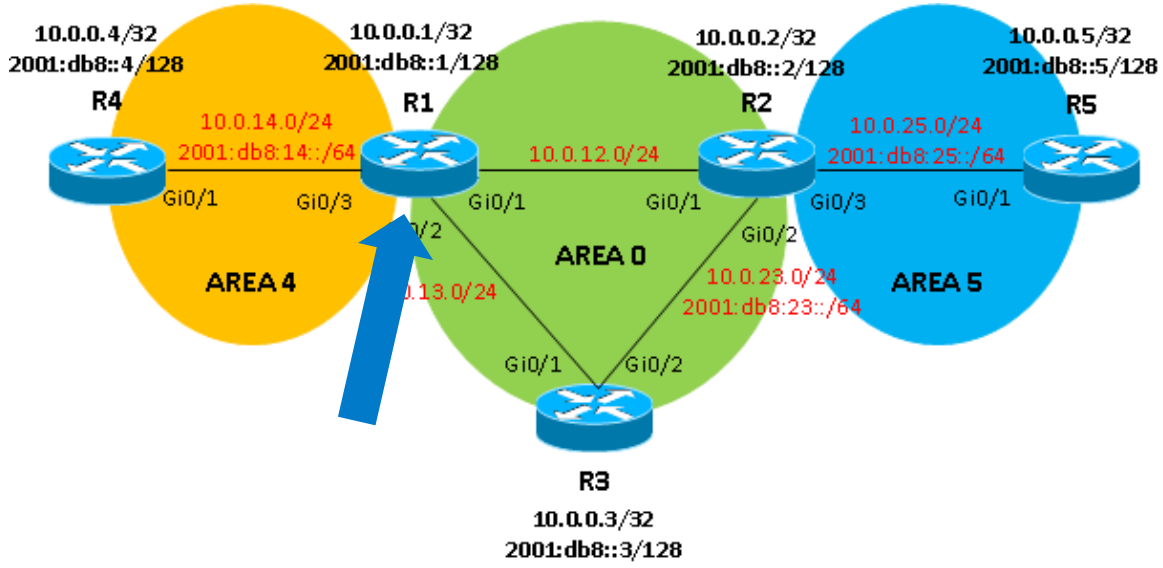

© 2014 Cisco and/or its affiliates. All rights reserved.

IS-IS for IPv6

IS-IS for IPv6

Implementation	Two new TLVs: - IPv6 Reachability TLV-236 (0xEC): Describes network reachability (IPv6 routing prefix, metric information and option bits) - IPv6 Interface Address TLV-232 (0xE8): Contains 128 bit address. Hello PDUs, must contain the link-local address but for LSP, must only contain the non link-local address A new Network Layer Protocol Identifier (NLPID): Allows IS-IS routers to advertise IPv6 prefix payload using 0x8E value (IPv4: 0xCC). Carried in Protocols Supported TLV (0x81).
Operational Considerations	Single Topology (default for IOS) - potentially beneficial in saving resources (same topology and same SPF) Multi Topology (RFC 5120) (default for IOS-XR) - Independent IPv4 and IPv6 topologies (MT ID 0,2), independent interface metrics. Wide metrics. New TLVs! Transition mode available - both types of TLVs are advertised
Notes	Standardization: Single Topology: http://tools.ietf.org/html/rfc5308 Multi Topology: http://tools.ietf.org/html/rfc5120 Evolution – IS-IS Multi-Instance: http://tools.ietf.org/html/rfc6822

© 2014 Cisco and/or its affiliates. All rights reserved.

IS-IS for IPv4 and IPv6 Summary

	IS-IS for IPv4 (RFC 1195)	IS-IS for IPv6 (RFC 5308)
Type	Link state	Link state

For Your Reference

IS-IS for IPv4

```
interface Loopback0
 ip address 10.10.0.1 255.255.255.255
 ip router isis CISCO
 isis circuit-type level-1
 isis metric 10000
!
interface Ethernet0/0
 ip address 10.0.0.1 255.255.255.0
 ip router isis CISCO
 isis circuit-type level-1
 isis metric 10000
!
router isis CISCO
 net 49.0001.1111.1111.1111.00
 metric-style wide
 log-adjacency-changes all
```

IS-IS for IPv6 – Single Topology

```
ipv6 unicast-routing
!
interface Loopback0
 ipv6 address 2001:DB8:1000::1/128
 ipv6 router isis CISCO
 isis circuit-type level-1
 isis ipv6 metric 10000
!
interface Ethernet0/0
 ipv6 address 2001:DB8::1/64
 ipv6 router isis CISCO
 isis circuit-type level-1
 isis ipv6 metric 10000
!
router isis CISCO
 net 49.0001.1111.1111.1111.00
 metric-style wide
 log-adjacency-changes all
```

IS-IS for IPv6 – Multi Topology

```
ipv6 unicast-routing
!
interface Ethernet0/0
 ipv6 address 2001:DB8::1/64
 ipv6 router isis CISCO
 isis circuit-type level-1
 isis ipv6 metric 10000
!
router isis CISCO
 net 49.0001.2222.2222.2222.00
 metric-style wide
 log-adjacency-changes all
 !
 address-family ipv6
  multi-topology
 exit-address-family
```


© 2014 Cisco and/or its affiliates. All rights reserved.

Single Topology IS-IS for IPv6

- IS-IS uses the same SPT topology for both IPv4 and IPv6
- Not suitable for an existing IPv4 IS-IS network where customer wants to turn on scattered IPv6 support
- Suitable for a network wide deployment where congruent IPv4/IPv6 topologies are required
  - Minimizes hardware resources impact (CPU + memory)
  - Provides consistent network convergence behaviour between for IPv4 and IPv6 topologies
- IPv4 and IPv6 topologies MUST match exactly
  - If not, black-holing of traffic can occur
  - In case of mismatching configurations on both ends of a Level-2 adjacency -> adjacency will establish (but black-holing can still occur)
  - In case of mismatching configurations on both ends of a Level-1 adjacency → adjacency will fail
- Therefore, during migration (from an IPv4-only to a Single Topology IPv4/IPv6 network) adjacency checking should be disabled – "no adjacency-check"

© 2014 Cisco and/or its affiliates. All rights reserved.

Single Topology IS-IS for IPv6 Example

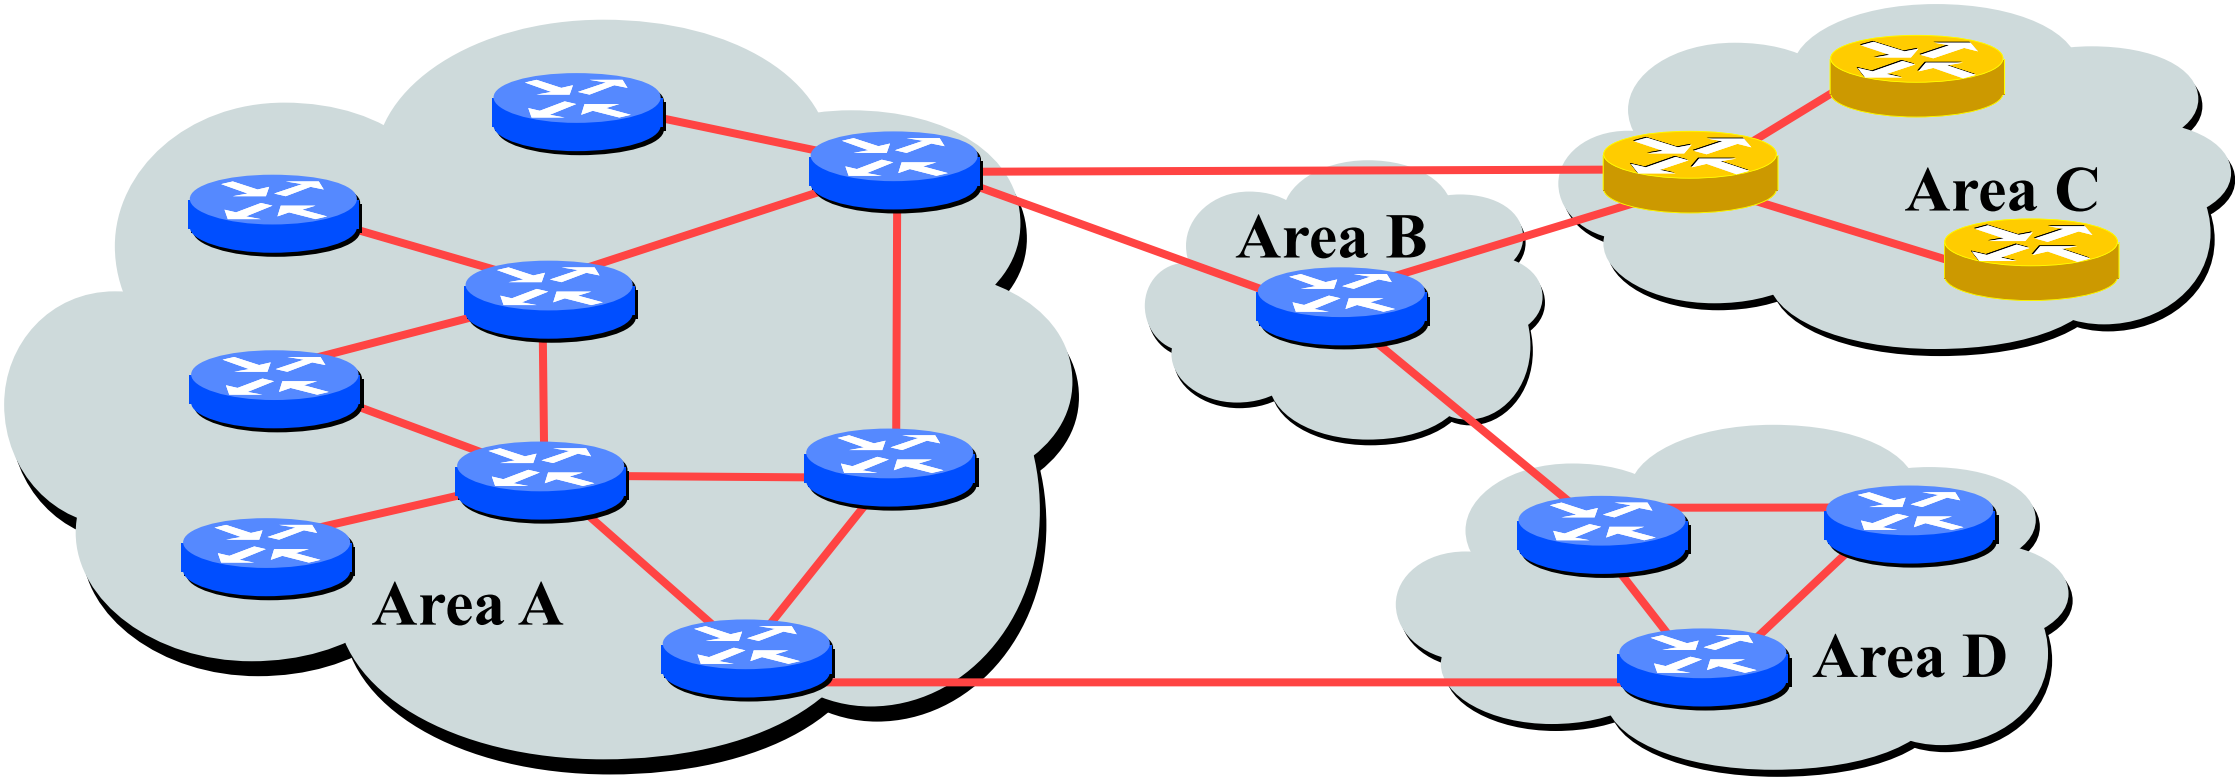

IPv4-IPv6 enables router

IPv4-only enabled router

© 2014 Cisco and/or its affiliates. All rights reserved.

Single Topology IS-IS Restriction

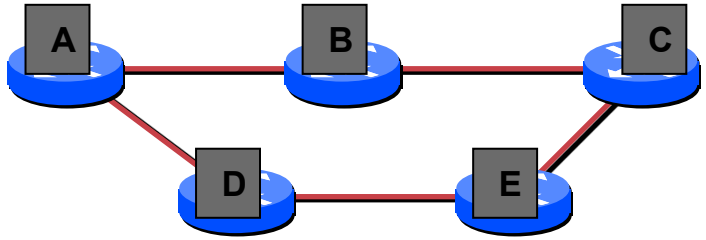

Physical Topology

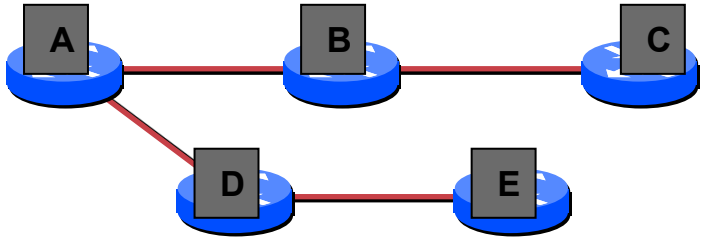

Shortest Path Tree

- From A's perspective, C is only reachable through B
  - There is no path from E to C (from A's perspective)
- All protocols carried by IS-IS have to agree on the same SPT
  - Not possible to distribute traffic across the domain
  - All links need to understand all protocols

© 2014 Cisco and/or its affiliates. All rights reserved.

Multi Topology IS-IS for IPv6 Solution

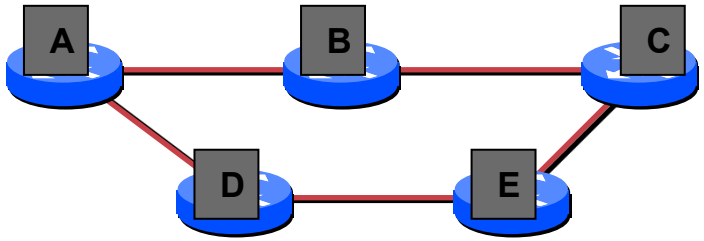

Physical Topology

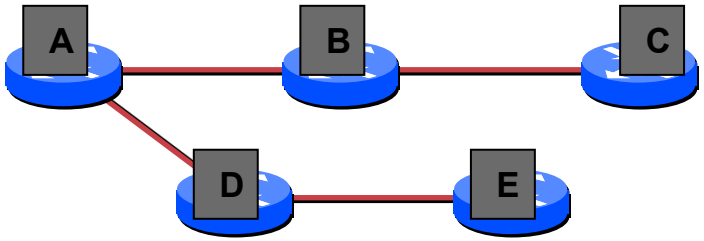

IPv4 Shortest Path Tree

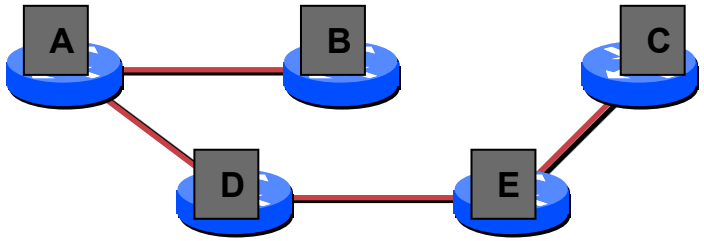

IPv6 Shortest Path Tree

- Ability to distribute traffic across all links
  - Separate traffic per address family
- Allows the deployment of non-congruent IPv4/IPv6 topologies

© 2014 Cisco and/or its affiliates. All rights reserved.

Multi Topology IS-IS

- IS-IS uses distinct SPTs for the IPv4 and IPv6 topologies
- Allows the deployment of non-congruent IPv4/IPv6 topologies
  - Allows for scattered / limited IPv6 support in an existing IPv4 IS-IS network
  - When congruent IPv4/IPv6 topologies are not possible / not desirable
  - Hardware resources impact (CPU + memory) is higher compared to Single Topology IS-IS (mostly not an issue with modern distributed platforms with powerful RPs)
  - Allows also for tuning of per-SPT topology tuning (e.g. network convergence timers)
- At adjacency establishment, peers will agree on topologies (SPTs) to be supported
  - Topologies identifiers are exchanged in IIH packets
  - L1 / L2 boundaries must be identical for all topologies
- Transition mode "multi-topology transition" is available to support both sets of TLVs
- Defaults!!!
  - IOS: Single Topology
  - IOS-XR: Multi Topology

© 2014 Cisco and/or its affiliates. All rights reserved.

Multi Topology IS-IS for IPv6 Example

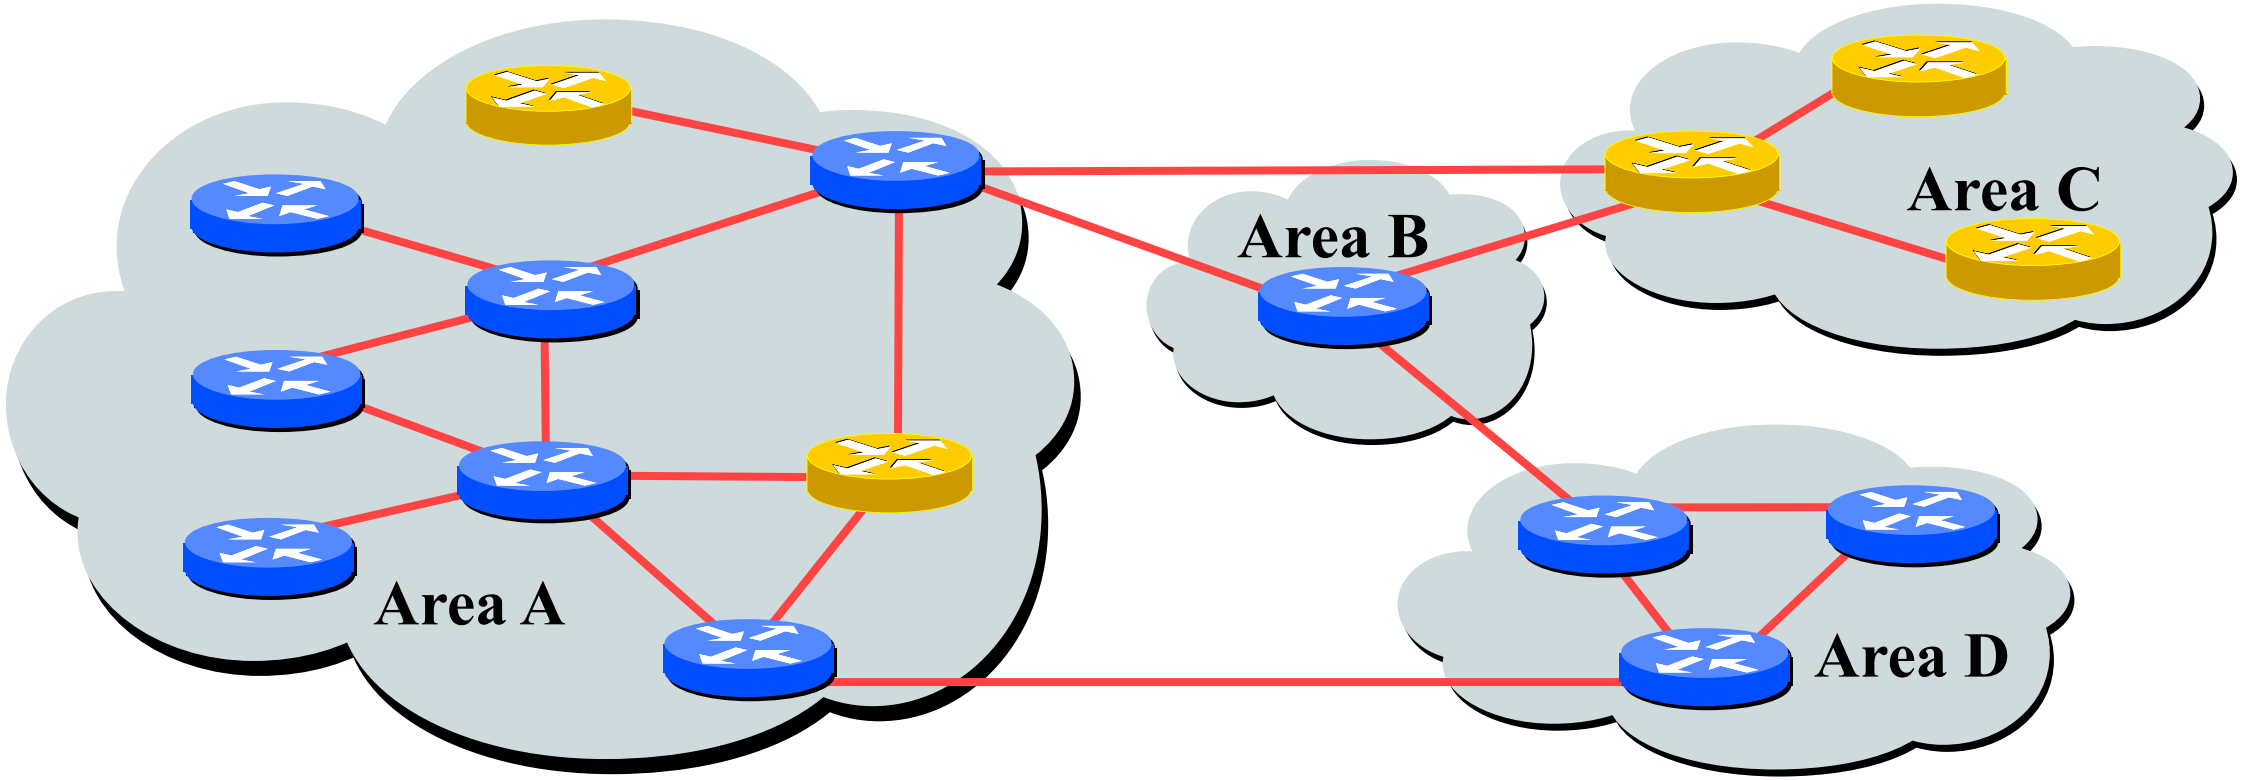

IPv4-IPv6 enabled router

IPv4-only enabled router

Multi Topology IS-IS will create two topologies inside each area for IPv4 and IPv6.
IPv4-only routers will be excluded from the IPv6 topology

© 2014 Cisco and/or its affiliates. All rights reserved.

Multi Topology IS-IS – New TLVs

- New TLVs used in IIHs and LSPs
  - TLV-229: Multi Topology TLV → contains Multi Topology Identifiers (MT IDs)
  - TLV-222: Multi Topology Intermediate System TLV
  - TLV-235: Multi Topology Reachable IPv4 prefixes TLV
  - TLV-237: Multi Topology Reachable IPv6 prefixes
- The Multi Topology Identifier (MT ID) describes the address family of the topology
- Reserved MT ID values are:
  - MT ID #0: Equivalent to the “standard” topology.
  - MT ID #1: Reserved for IPv4 in-band management purposes.
  - MT ID #2: Reserved for IPv6 routing topology.
  - MT ID #3: Reserved for IPv4 multicast routing topology.
  - MT ID #4: Reserved for IPv6 multicast routing topology.
  - MT ID #5: Reserved for IPv6 in-band management purposes.
  - MT ID #6-#3995: Reserved for IETF consensus.
  - MT ID #3996-#4095: Reserved for development, experimental and proprietary features.

For Your Reference

© 2014 Cisco and/or its affiliates. All rights reserved.

IS-IS Sample Topology

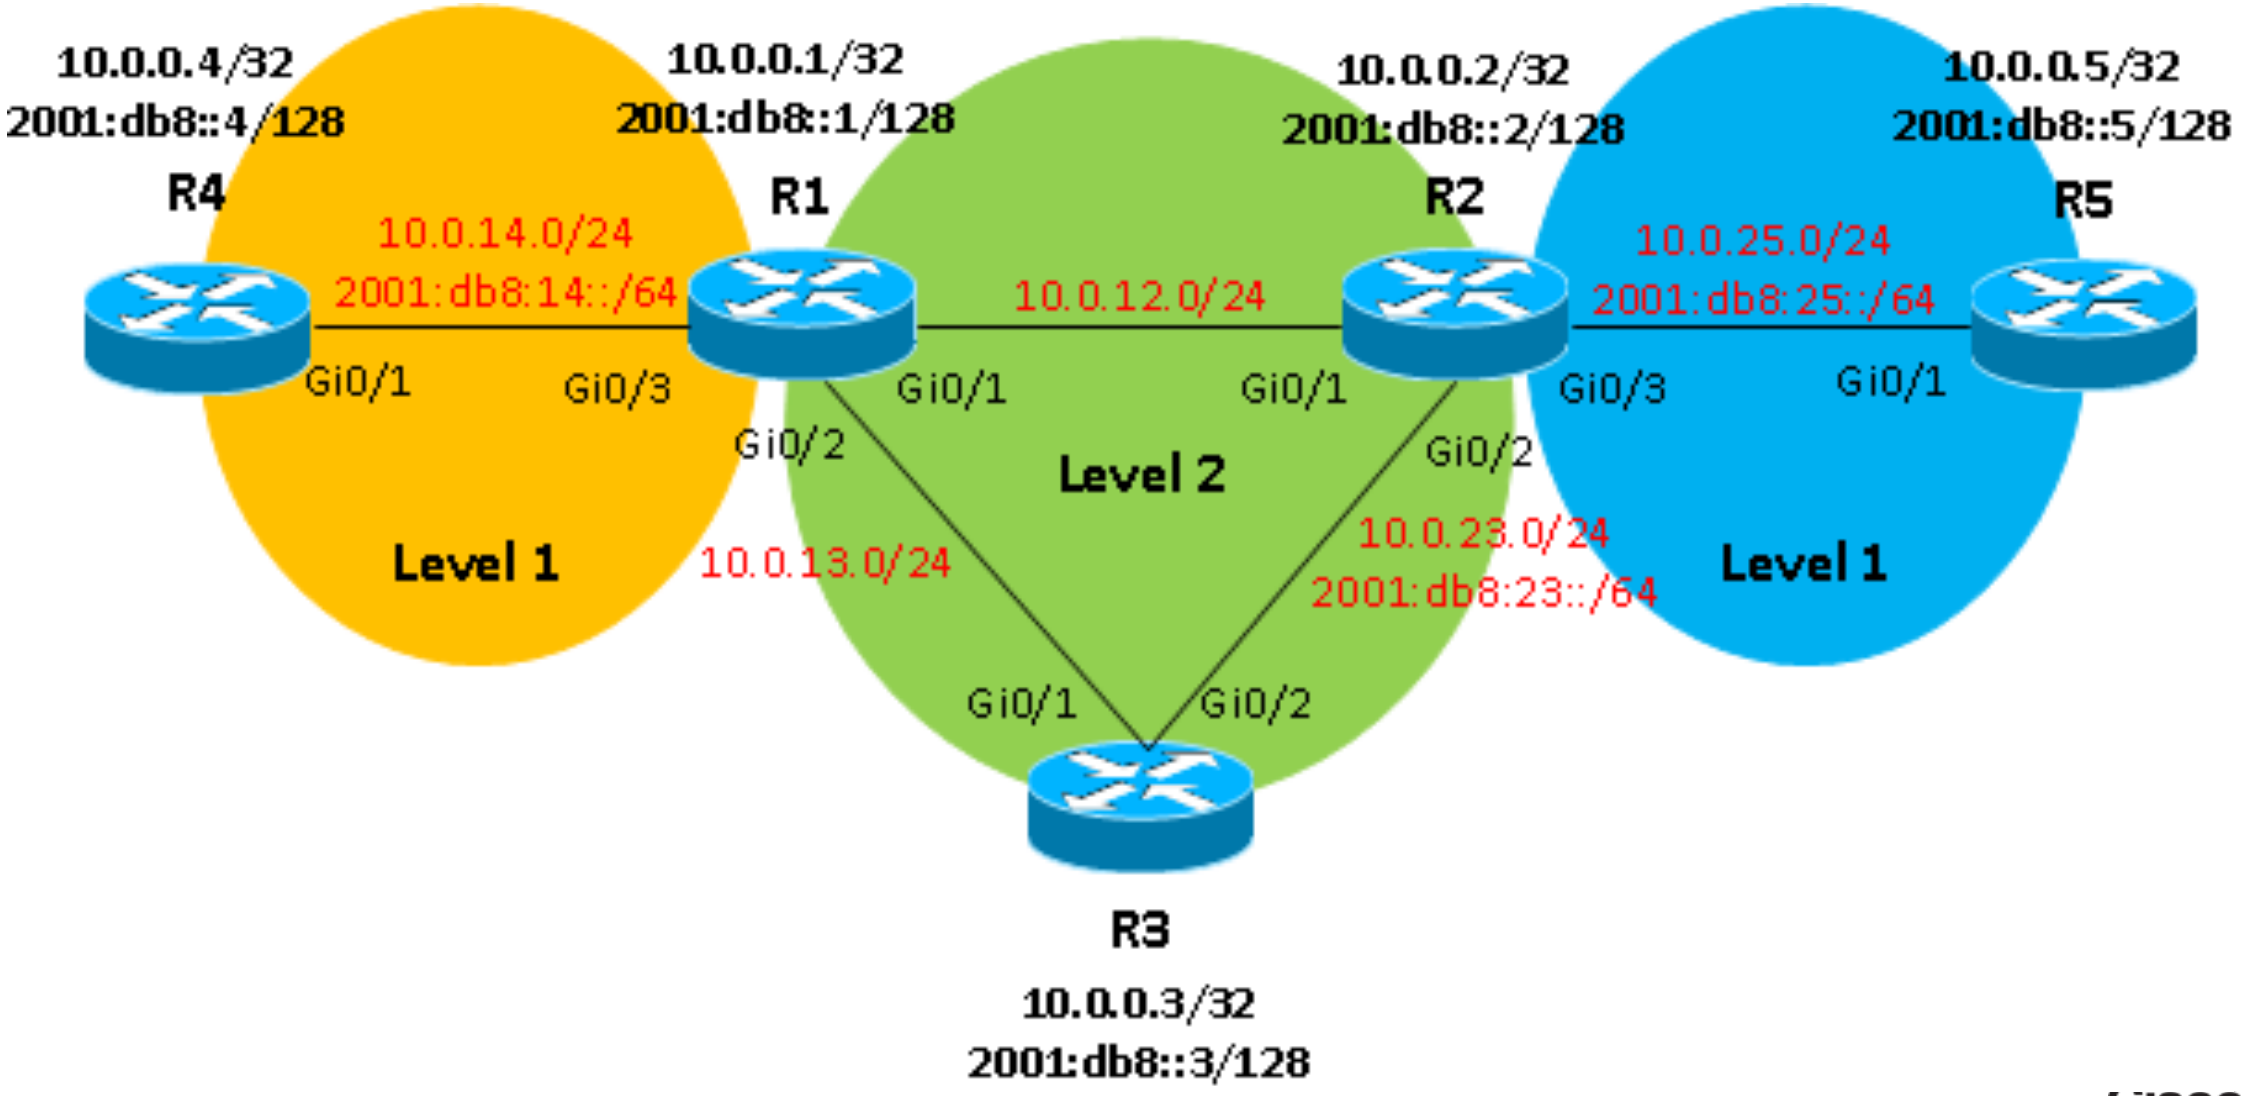

© 2014 Cisco and/or its affiliates. All rights reserved.

IS-IS Configuration (Single-Topology)

```
!
ipv6 unicast-routing
!
interface Loopback0
 ip address 10.0.0.1 255.255.255.255
 ip router isis CiscoLive
 ipv6 address 2001:DB8::1/128
 ipv6 router isis CiscoLive
 isis circuit-type level-2-only
!
interface GigabitEthernet0/1
 ip address 10.0.12.1 255.255.255.0
 ip router isis CiscoLive
 ipv6 enable
 ipv6 router isis CiscoLive
 isis circuit-type level-2-only
 isis network point-to-point
!
interface GigabitEthernet0/2
 ip address 10.0.13.1 255.255.255.0
 ip router isis CiscoLive
 ipv6 enable
 ipv6 router isis CiscoLive
 isis circuit-type level-2-only
 isis network point-to-point
!
interface GigabitEthernet0/3
 ip address 10.0.14.1 255.255.255.0
 ip router isis CiscoLive
 ipv6 enable
 ipv6 router isis CiscoLive
 isis circuit-type level-1
 isis network point-to-point
```

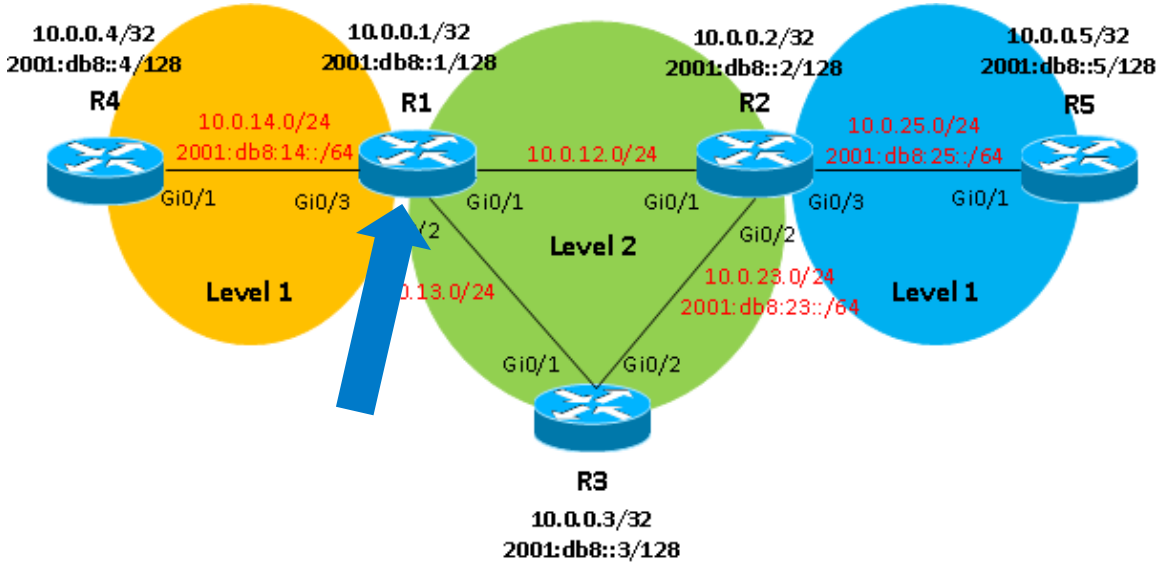


```
....
!
router isis CiscoLive
 net 49.0004.1111.1111.1111.00
 metric-style wide
 log-adjacency-changes all
!
```

© 2014 Cisco and/or its affiliates. All rights reserved.

IS-IS Interfaces

```
R1#show ipv6 interface gigabitEthernet 0/1
GigabitEthernet0/1 is up, line protocol is up
  IPv6 is enabled, link-local address is FE80::F816:3EFF:FECB:
18A5
  No V
  No g
  Join
    FF
    FF
    FF
  MTU
  ICMP
  ICMP
  ICMP
  ND D
  ND r
  ND a
  ND a
  ND r
  ND r
  ND a
  Host
R1#
```

```
R1#show clns interface gigabitEthernet 0/1
GigabitEthernet0/1 is up, line protocol is up
  Che
  ERP
  CLN
  CLN
  DEC
  Nex
  Rou
    C
    I
    N
    I
    I
    N
    N
    i
R1#
```

```
R3#show clns interface gigabitEthernet 0/2
GigabitE
  Checks
  ERPDUs
  CLNS f
  CLNS S
  DEC co
  Next E
  Routin
    Circ
    Inte
    Leve
    DR I
    Leve
    Numb
    Next
R3#
```

```
R1#show clns protocol

IS-IS Router: CiscoLive
  System Id: 1111.1111.1111.00  IS-Type: level-1-2
  Manual area address(es):
        49.0004
  Routing for area address(es):
        49.0004
  Interfaces supported by IS-IS:
        Loopback0 - IP - IPv6
        GigabitEthernet0/3 - IPv6
        GigabitEthernet0/2 - IP - IPv6
        GigabitEthernet0/1 - IP - IPv6
  Redistribute:
    static (on by default)
  Distance for L2 CLNS routes: 110
  RRR level: none
  Generate narrow metrics: none
  Accept narrow metrics:   none
  Generate wide metrics:   level-1-2
  Accept wide metrics:     level-1-2
R1#
```

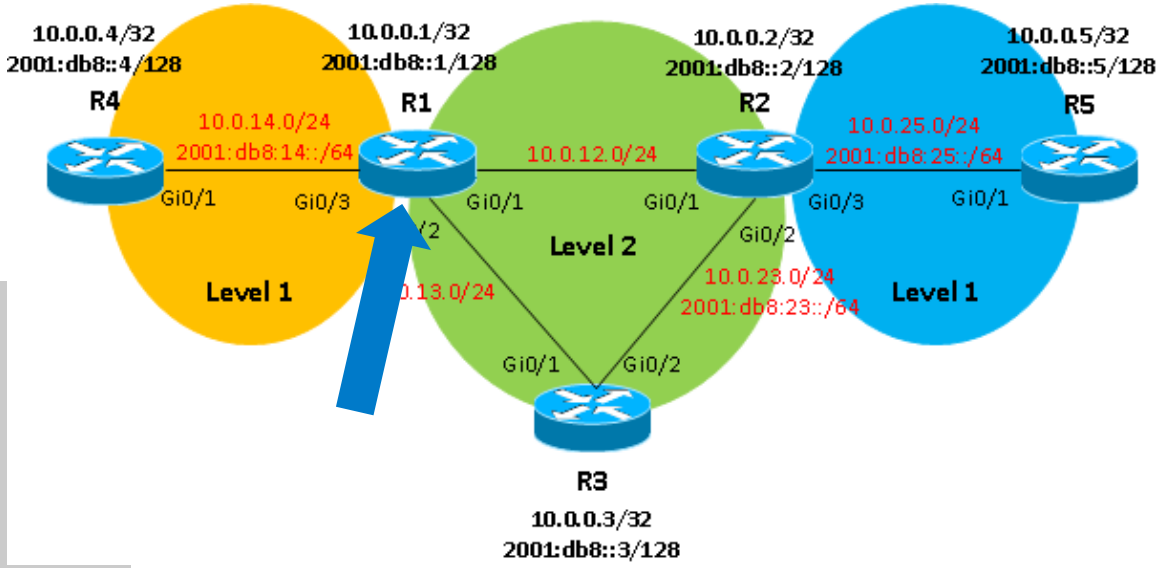

© 2014 Cisco and/or its affiliates. All rights reserved.

IS-IS Neighbours

```
R1#show isis neighbors detail

Tag CiscoLive:
System Id       Type Interface    IP Address       State Holdtime Circuit
Id
R2              L2   Gi0/1        10.0.12.2        UP    20       01
  Area Address(es): 49.0005
  SNPA: fa16.3e46.08b0
  IPv6 Address(es): FE80::F816:3EFF:FE46:8B0
  State Changed: 00:07:06
  Format: Phase V
  Remote TID: 0
  Local TID:  0
  Interface name: GigabitEthernet0/1
R3              L2   Gi0/2        10.0.13.3        UP    21       01
  Area Address(es): 49.0001
  SNPA: fa16.3ef7.8f20
  IPv6 Address(es): FE80::F816:3EFF:FEF7:8F20
  State Changed: 00:23:04
  Format: Phase V
  Remote TID: 0
  Local TID:  0
  Interface name: GigabitEthernet0/2
R4              L1   Gi0/3                         UP    :
  Area Address(es): 49.0004
  SNPA: fa16.3ecf.eed2
  IPv6 Address(es): FE80::F816:3EFF:FECF:EED2
  State Changed: 00:22:45
  Format: Phase V
  Remote TID: 0
  Local TID:  0
  Interface name: GigabitEthernet0/3
R1#
```

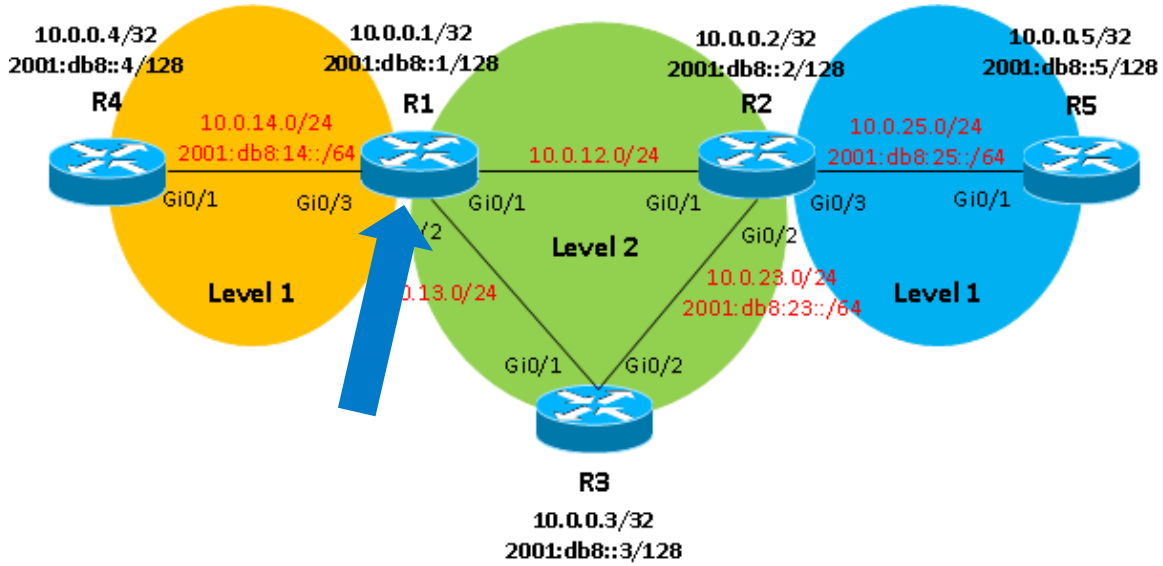


```
R1#show clns neighbors gigabitEthernet 0/1 detail

Tag CiscoLive:
System Id       Interface    SNPA                State   Holdtime   Type
Protocol
R2              Gi0/1        fa16.3e46.08b0      Up      20         L2    IS-IS
  Area Address(es): 49.0005
  IP Address(es):  10.0.12.2*
  IPv6 Address(es): FE80::F816:3EFF:FE46:8B0
  Uptime: 00:08:07
  NSF capable
  Interface name: GigabitEthernet0/1
R1#
```

IS-IS RIB

```
R1#show ipv6 route isis
IPv6 Routing Table - default - 8 entries
Codes: C - Connected, L - Local, S - Static, U - Per-user Static route
       B - BGP, HA - Home Agent, MR - Mobile Router, R - RIP
       H - NHRP, I1 - ISIS L1, I2 - ISIS L2, IA - ISIS interarea
       IS - ISIS summary, D - EIGRP, EX - EIGRP external, NM - NEMO
       ND - ND Default, NDp - ND Prefix, DCE - Destination, NDr -
Redirect
       O - OSPF Intra, OI - OSPF Inter, OE1 - OSPF ext 1, OE2 - OSPF ext
2
       ON1 - OSPF NSSA ext 1, ON2 - OSPF NSSA ext 2, ls - LISP site
       ld - LISP dyn-EID, a - Application
I2  2001:DB8::2/128 [115/20]
     via FE80::F816:3EFF:FE46:8B0, GigabitEthernet0/1
I2  2001:DB8::3/128 [115/20]
```

```
R1#show isis ipv6 rib
IS-IS IPv6 process CiscoLive, local RIB
* 2001:DB8::2/128
    via FE80::F816:3EFF:FE46:8B0/GigabitEthernet0/1, type L2  metric 20 tag 0 LSP [4/20]
* 2001:DB8::3/128
    via FE80::F816:3EFF:FEF7:8F20/GigabitEthernet0/2, type L2  metric 20 tag 0 LSP [6/1A]
* 2001:DB8::4/128
    via FE80::F816:3EFF:FECF:EED2/GigabitEthernet0/3, type L1  metric 20 tag 0 LSP [8/B]
* 2001:DB8::5/128
    via FE80::F816:3EFF:FE46:8B0/GigabitEthernet0/1, type L2  metric 30 tag 0 LSP [4/20]
* 2001:DB8:23::/64
    via FE80::F816:3EFF:FEF7:8F20/GigabitEthernet0/2, type L2  metric 20 tag 0 LSP [6/1A]
    via FE80::F816:3EFF:FE46:8B0/GigabitEthernet0/1, type L2  metric 20 tag 0 LSP [4/20]
* 2001:DB8:25::/64
    via FE80::F816:3EFF:FE46:8B0/GigabitEthernet0/1, type L2  metric 20 tag 0 LSP [4/20]
R1#
```

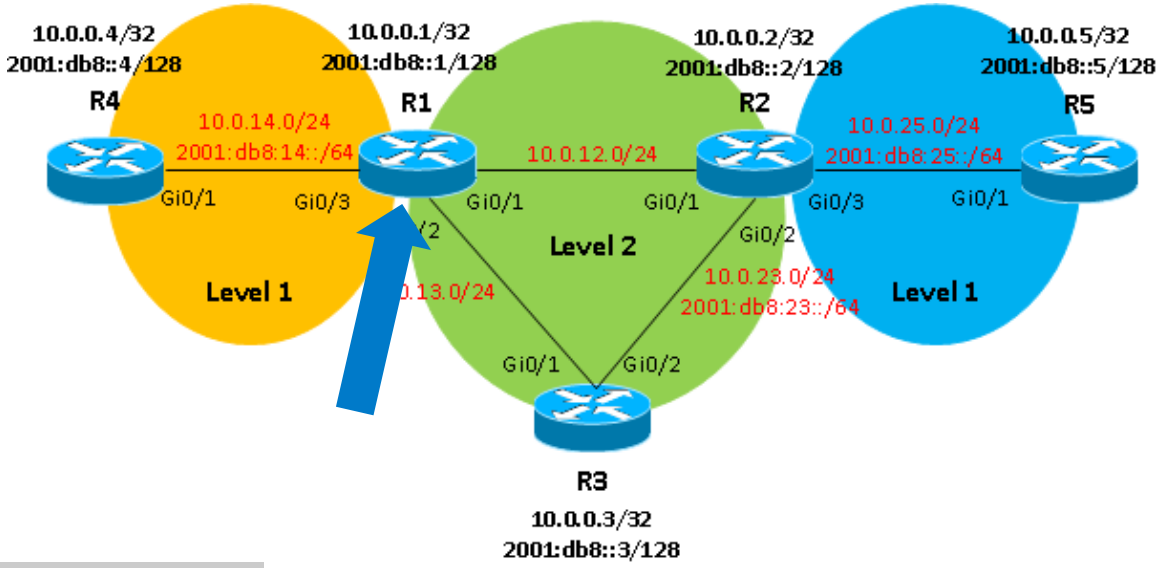

© 2014 Cisco and/or its affiliates. All rights reserved.
Cisco Public

IS-IS LSDB

```
R1#show isis database detail

Tag CiscoLive:
....
IS-IS Level-2 Link State Database:
LSPID                 LSP Seq Num  LSP Checksum  LSP Holdtime      ATT/P/
OL
R1.00-00            * 0x00000015   0x591F        723               0/0/0
  Area Address: 49.0004
  NLPID:        0xCC 0x8E
  Hostname: R1
  IP Address:
  Me
  Me
  Me
  IP
  Me
  Me
  Me
  Me
...
```

```
....
R2.00-00              0x00000018   0x3EC3        929               0/0/
  Area Address: 49.0005
  NLPID:        0xCC 0x8E
  Hostname: R2
  IP Address:   10.0.0.2
  Metric: 10         IP 10.0.0.2/32
  Metric: 10         IP 10.0.12.0/24
  Metric: 10         IP 10.0.23.0/24
  IPv6 Address: 2001:DB8::2
  Metric: 10         IPv6 2001:DB8:23::/64
  Metric: 10         IPv6 2001:DB8::2/128
  Metric: 10         IS-Extended R2.02
  Metric: 10         IS-Extended R1.00
  Metric: 10         IPv6 2001:DB8:25::/64
  Metric: 20         IPv6-Interarea 2001:DB8::5/128
R2.02-00              0x00000006   0x24D7        1088              0/0/0
  Metric: 0          IS-Extended R2.00
  Metric: 0          IS-Extended R3.00
....
R1#
```

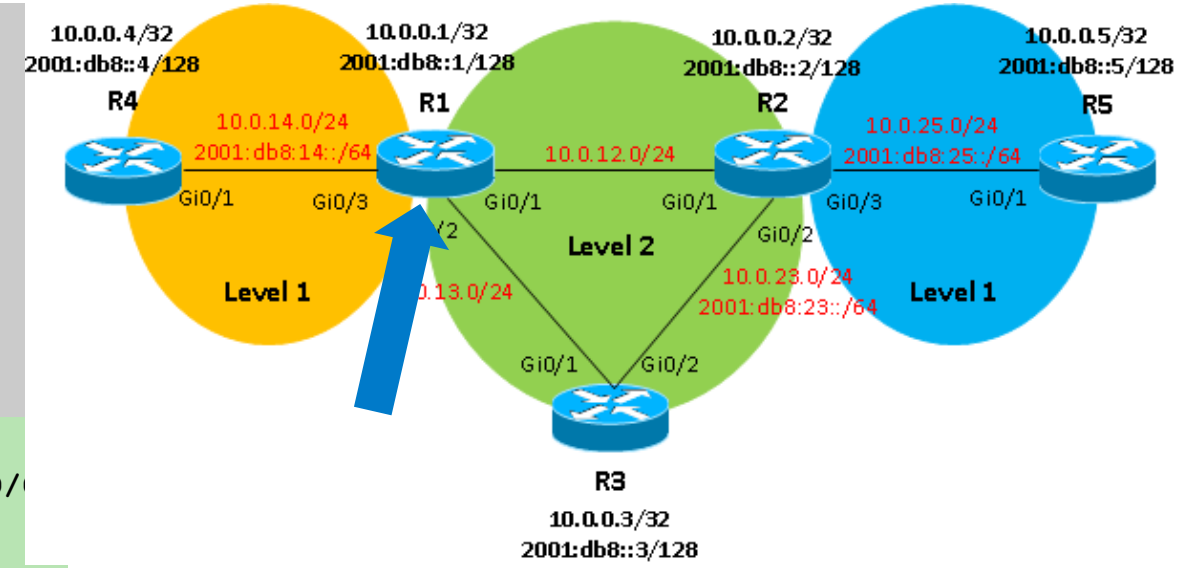

IS-IS Multi-Topology Configuration

```
!
router isis CiscoLive
 net 49.0004.1111.1111.1111.00
 metric-style wide
 log-adjacency-changes all
 !
 address-family ipv6
  multi-topology
 exit-address-family
!
```

OR

```
!
router isis CiscoLive
 net 49.0004.1111.1111.1111.00
 metric-style wide
 log-adjacency-changes all
 !
 address-family ipv6
  multi-topology transition
 exit-address-family
!
```

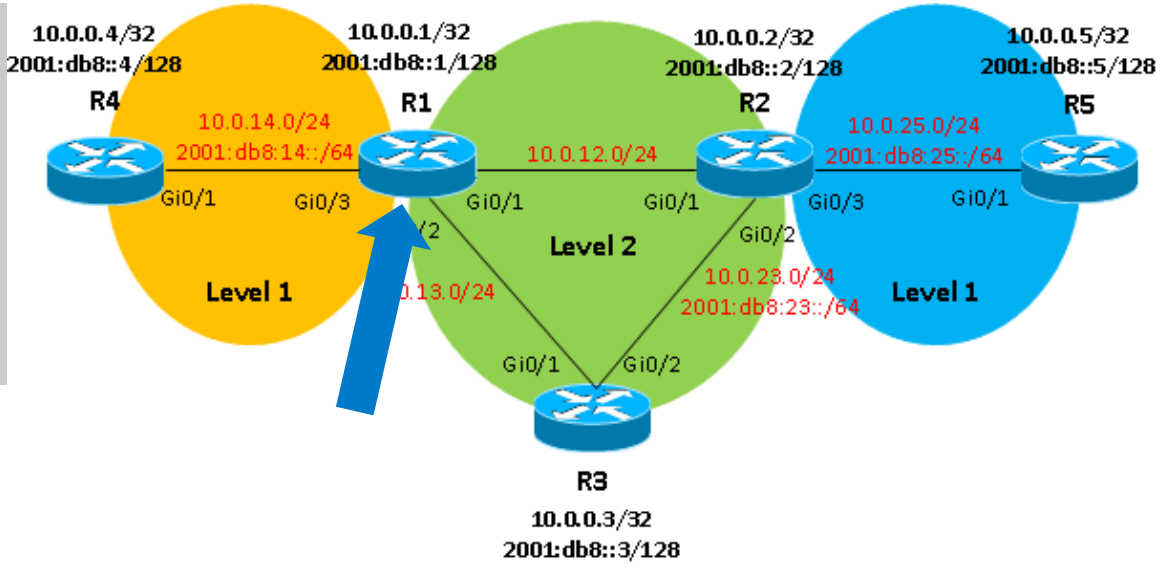

© 2014 Cisco and/or its affiliates. All rights reserved.

IS-IS Neighbours

```
R1#show isis neighbors detail

R1#show clns neighbors detail
Tag CiscoLive:
System Id      Interface   SNPA            State  Holdtime  Type
Protocol
R2             Gi0/1       fa16.3e46.08b0  Up     22        L2   M-
ISIS
  Area Address(es): 49.0005
  IP Address(es):  10.0.12.2*
  IPv6 Address(es): FE80::F816:3EFF:FE46:8B0
  Uptime: 13:21:40
  NSF capable
  Topology: IPv4, IPv6
  Interface name: GigabitEthernet0/1
R3             Gi0/2       fa16.3ef7.8f20  Up     26        L2   M-
ISIS
  Area Address(es): 49.0001
  IP Address(es):  10.0.13.3*
  IPv6 Address(es): FE80::F816:3EFF:FEF7:8F20
  Uptime: 1d12h
  NSF capable
  Topology: IPv4, IPv6
  Interface name: GigabitEthernet0/2
R4             Gi0/3       fa16.3ecf.eed2  Up     25        L1   M-
ISIS
  Area Address(es): 49.0004
  IP Address(es):  10.0.14.4*
  IPv6 Address(es): FE80::F816:3EFF:FECF:EED2
  Uptime: 12:23:50
System Id      Interface   SNPA            State  Holdtime  Type
Protocol
  NSF capable
  Topology: IPv4, IPv6
```

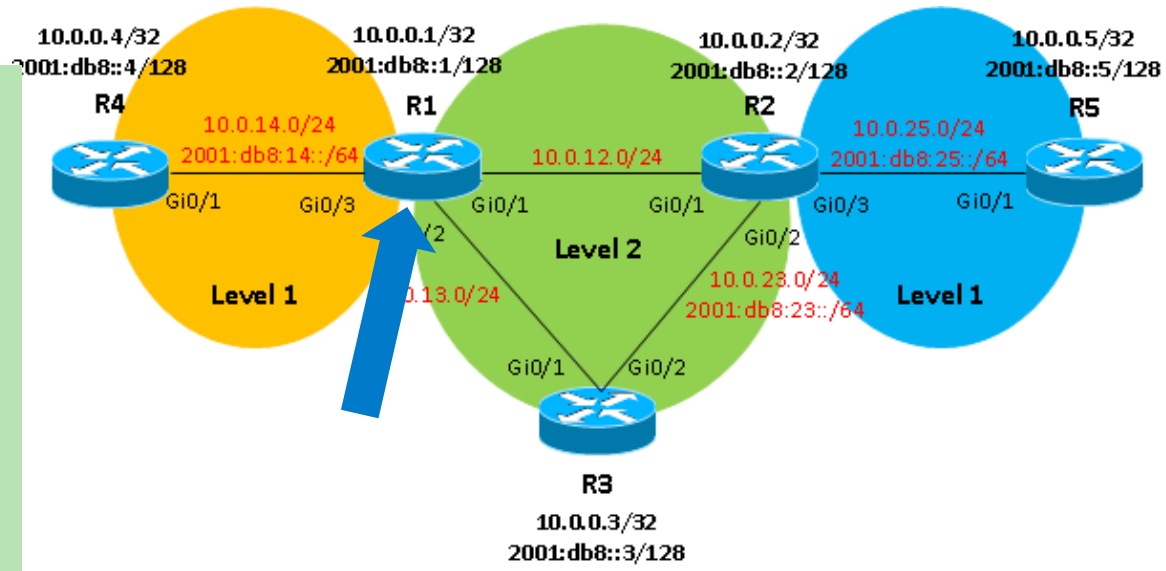

IS-IS Multi-Topology Transition LSDB

```
R1#show isis database detail

Tag CiscoLive:
....
IS-IS Level-2 Link State Database:
LSPID                 LSP Seq Num  LSP Checksum  LSP Holdtime
P/OL
R1.00-00            * 0x0000008B   0xDE92        892
  Area Address: 49.0004
  Topology:     IPv4 (0x0)
                IPv6 (0x2)
  NLPID:        0xCC 0x8E
  Hostname: R1
  IP Address:   10.0.0.1
  Metric: 10         IP 10.0.0.1/32
  Metric: 10         IP 10.0.12.0/24
  Metric: 10         IP 10.0.13.0/24
  IPv6 Address: 2001:DB8::1
  Metric: 10         IPv6 2001:DB8::1/128
  Metric: 10         IPv6 (MT-IPv6) 2001:DB8::1/128
  Metric: 10         IS-Extended R2.00
  Metric: 10         IS-Extended R3.00
  Metric: 10         IS (MT-IPv6) R2.00
  Metric: 10         IS (MT-IPv6) R3.00
  Metric: 20         IP 10.0.0.4/32
  Metric: 10         IP 10.0.14.0/24
  Metric: 20         IPv6-Interarea 2001:DB8::4/128
  Metric: 20         IPv6-Interarea (MT-IPv6) 2001:DB8::4/128
....
```

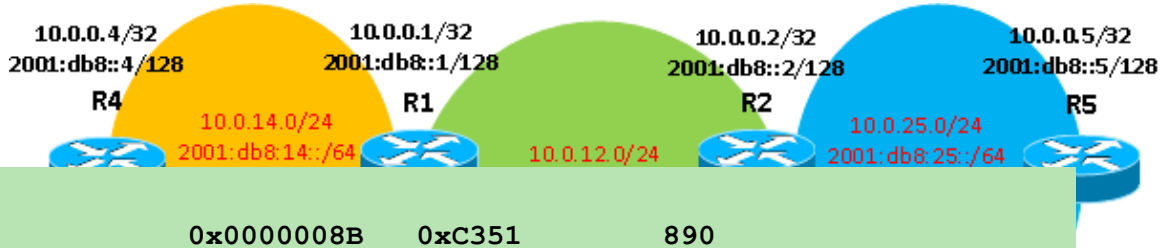


```
....
R2.00-00              0x0000008B   0xC351        890
0/0/0
  Area Address: 49.0005
  Topology:     IPv4 (0x0)
                IPv6 (0x2)
  NLPID:        0xCC 0x8E
  Hostname: R2
  IP Address:   10.0.0.2
  Metric: 10         IP 10.0.0.2/32
  Metric: 10         IP 10.0.12.0/24
  Metric: 10         IP 10.0.23.0/24
  IPv6 Address: 2001:DB8::2
  Metric: 10         IPv6 2001:DB8:23::/64
  Metric: 10         IPv6 2001:DB8::2/128
  Metric: 10         IPv6 (MT-IPv6) 2001:DB8:23::/64
  Metric: 10         IPv6 (MT-IPv6) 2001:DB8::2/128
  Metric: 10         IS-Extended R2.02
  Metric: 10         IS-Extended R1.00
  Metric: 10         IS (MT-IPv6) R2.02
  Metric: 10         IS (MT-IPv6) R1.00
  Metric: 10         IPv6 2001:DB8:25::/64
  Metric: 20         IPv6-Interarea 2001:DB8::5/128
  Metric: 10         IPv6 (MT-IPv6) 2001:DB8:25::/64
  Metric: 20         IPv6-Interarea (MT-IPv6) 2001:DB8::5/128
R2.02-00              0x00000074   0x4746        910
0/0/0
  Metric: 0          IS-Extended R2.00
  Metric: 0          IS-Extended R3.00
```


© 2014 Cisco and/or its affili

IS-IS Multi-Topology LSDB

```
R1#show isis database detail

Tag CiscoLive:
IS-IS Level-2 Link State Database:
LSPID                 LSP Seq Num  LSP Checksum  LSP Holdtime      ATT/P/
OL
R1.00-00            * 0x0000008D   0xC81F        454               0/0/0
  Area Address: 49.0004
  Top
  NLI
  Hos
  IP
  Met
  Met
  Met
  IPv
  Met
  Met
  Met
  Met
  Met
  Met
....
```

```
....
R2.00-00              0x0000008F   0xC78A        953               0/0/0
  Area Address: 49.0005
  Topology:     IPv4 (0x0)
                IPv6 (0x2)
  NLPID:        0xCC 0x8E
  Hostname: R2
  IP Address:   10.0.0.2
  Metric: 10         IP 10.0.0.2/32
  Metric: 10         IP 10.0.12.0/24
  Metric: 10         IP 10.0.23.0/24
  IPv6 Address: 2001:DB8::2
  Metric: 10         IPv6 (MT-IPv6) 2001:DB8:23::/64
  Metric: 10         IPv6 (MT-IPv6) 2001:DB8::2/128
  Metric: 10         IS-Extended R2.02
  Metric: 10         IS-Extended R1.00
  Metric: 10         IS (MT-IPv6) R2.02
  Metric: 10         IS (MT-IPv6) R1.00
  Metric: 10         IPv6 (MT-IPv6) 2001:DB8:25::/64
  Metric: 20         IPv6-Interarea (MT-IPv6) 2001:DB8::5/128
R2.02-00              0x00000075   0x4547        763               0/0/0
  Metric: 0          IS-Extended R2.00
  Metric: 0          IS-Extended R3.00
R1#
```

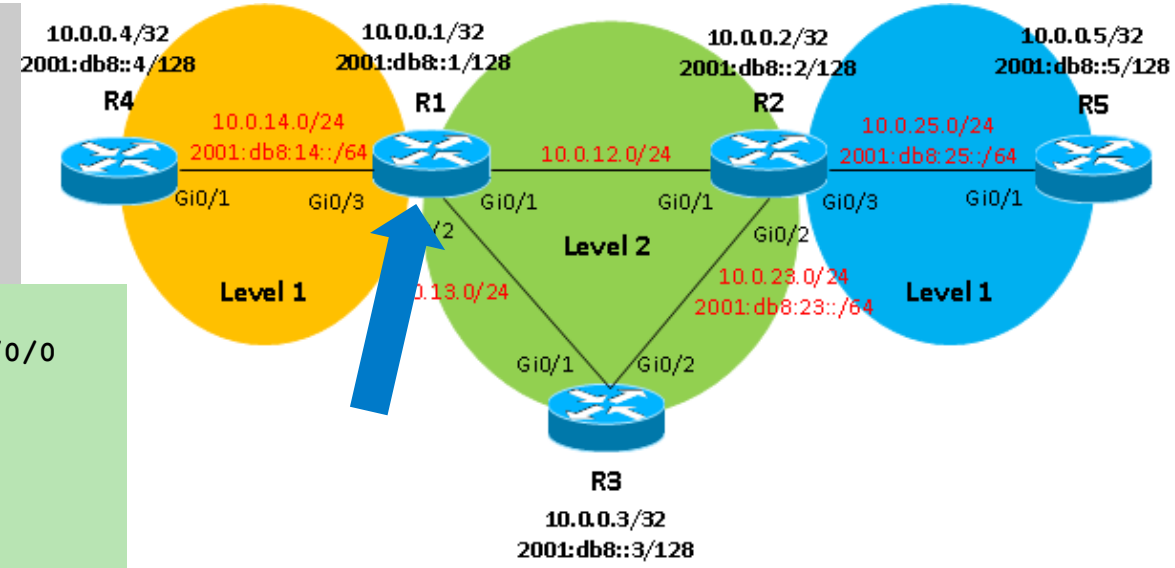

© 2014 Cisco and/or its affiliates. All rights reserved.
Cisco Public

BGP for IPv6

Multiprotocol Extensions for BGP-4 (RFC 4760)

- BGP-4 carries only 3 pieces of information which are IPv4 specific:
  - NLRI in the UPDATE message contains an IPv4 prefix
  - NEXT_HOP attribute in the UPDATE message contains a IPv4 address
  - BGP Identifier in the OPEN message & AGGREGATOR attribute
- To make BGP-4 available for other network layer protocols, RFC 4760 (obsoletes RFC 2858, 2283) defines multi-protocol extensions for BGP-4
  - Enables BGP-4 to carry information for other protocols e.g MPLS, IPv6
  - 2 new BGP-4 optional and non-transitive attributes:
  - Multiprotocol Reachable NLRI (MP_REACH_NLRI) – Type Code 14
  - Multiprotocol Unreachable NLRI (MP_UNREACH_NLRI) – Type Code 15
- MP_REACH_NLRI carries set of reachable destinations + next-hop information
- MP_UNREACH_NLRI carries set of unreachable destinations
- Both attributes contain the following information:
  - Address Family Identifies (AFI) + Subsequent Address Family Identifier (SAFI)
  - Protocol independent NEXT_HOP (MP_REACH_NLRI only)
  - Protocol independent NLRI / Withdrawn Routes

© 2014 Cisco and/or its affiliates. All rights reserved.

AFI / SAFI Values

- AFI
  - IPv4 NLRI: AFI=1
  - IPv6 NLRI: AFI=2
- SAFI
  - Unicast NLRI: SAFI=1
  - Multicast NLRI (for RPF check): SAFI=2
  - NLRI + MPLS label: SAFI=4
  - MPLS labelled VPN NLRI: SAFI=128
- For the full list of SAFI values: http://www.iana.org/assignments/safi-namespace/safi-namespace.xhtml

For Your Reference

© 2014 Cisco and/or its affiliates. All rights reserved.

BGP-4 Multiprotocol Extensions for IPv6 (RFC 2545)

- The next-hop field contains a global IPv6 address and potentially a link-local IPv6 address
- The value of the length of the next-hop field in the MP_REACH_NLRI attribute is set to 16 when only global is present and is set to 32 if the link-local is present as well
- A link-local address as a next-hop is only set if the BGP peer shares the subnet with both routers (advertising and advertised) – i.e. for eBGP
- In the event that only a link-local next-hop address is present, this needs to be changed to a global address for the iBGP update
  - Next-hop-self
  - Route-map
- TCP transport
  - BGP-4 runs on top of TCP
  - TCP can be setup on top of IPv4 or IPv6
- BGP Identifier "bgp router-id x.x.x.x"
  - Used within the OPEN message at session establishment between the BGP peers
  - Must be unique in network
  - In case no IPv4 is configured, an explicit "bgp router-id" needs to be configured

© 2014 Cisco and/or its affiliates. All rights reserved.

BGP for IPv4 / IPv6 Best Practices

- Use 2 BGP peers
  - IPv4 peer for IPv4 prefixes (address-family ipv4 unicast)
  - IPv6 peer for IPv6 prefixes (address-family ipv6 unicast)
- Allows independent control of IPv4 / IPv6 peering sessions
- # BGP sessions could pose scalability challenge
- IPv6 LL peer
  - Implicit protection for peer address – LL address is not reachable outside that link
  - Requires “hard-coded” LL address, requires “next-hop-self” or route-map for NH setting

© 2014 Cisco and/or its affiliates. All rights reserved.

BGP Sample Topology

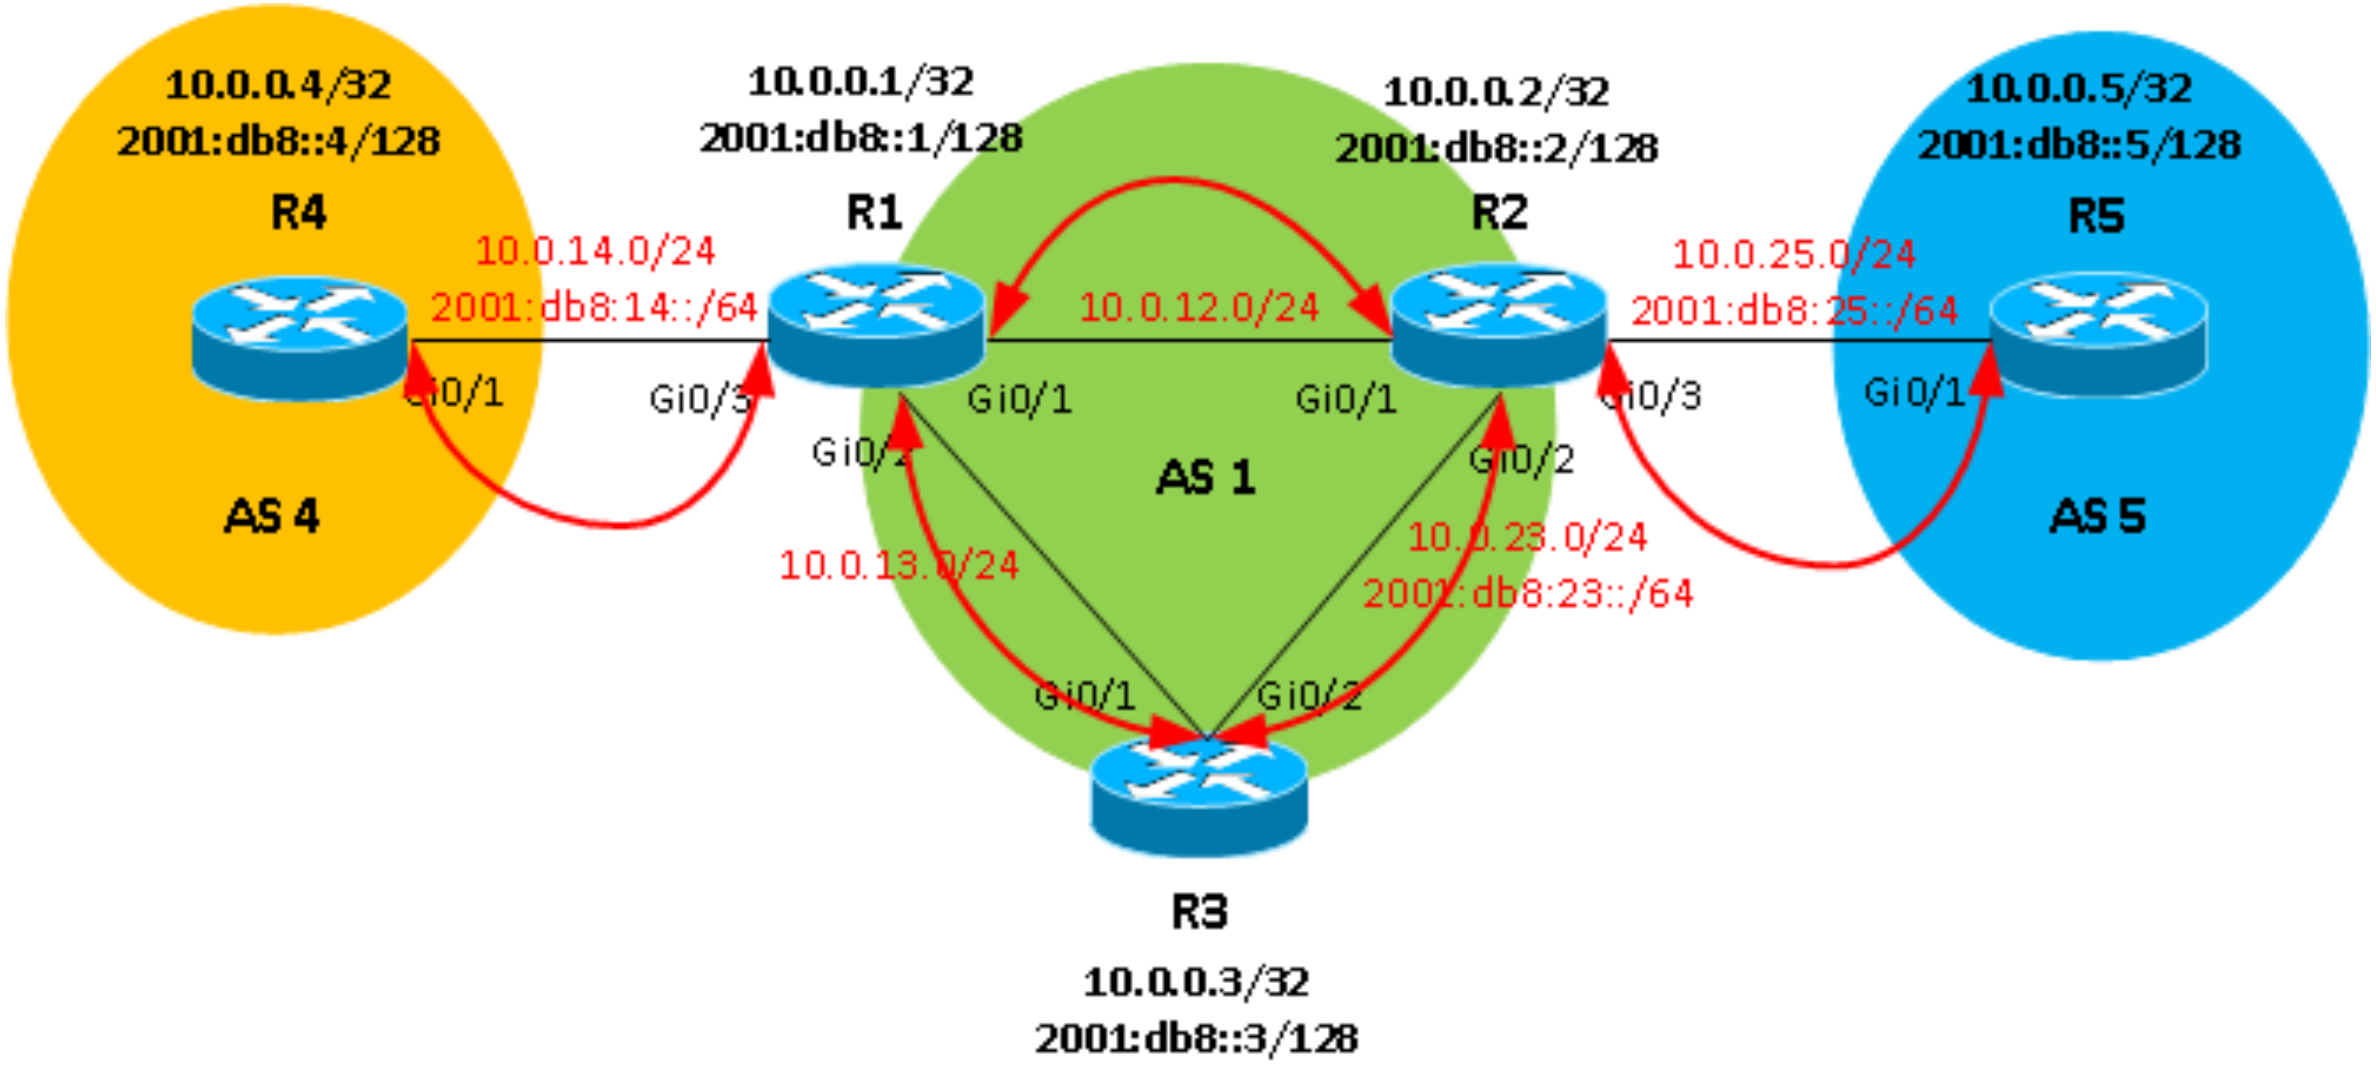

© 2014 Cisco and/or its affiliates. All rights reserved.

BGP Configuration

```
!
router bgp 1
 bgp router-id 10.0.0.1
 bgp log-neighbor-changes
 neighbor 10.0.0.2 remote-as 1
 neighbor 10.0.0.2 update-source Loopback0
 neighbor 10.0.0.3 remote-as 1
 neighbor 10.0.0.3 update-source Loopback0
 neighbor 10.0.14.4 remote-as 4
 neighbor 2001:DB8::2 remote-as 1
 neighbor 2001:DB8::2 update-source
Loopback0
 neighbor 2001:DB8::3 remote-as 1
 neighbor 2001:DB8::3 update-source
Loopback0
 neighbor 2001:DB8:14::4 remote-as 4
 !
 address-family ipv4
  neighbor 10.0.0.2 activate
  neighbor 10.0.0.2 next-hop-self
  neighbor 10.0.0.3 activate
  neighbor 10.0.0.3 next-hop-self
  neighbor 10.0.14.4 activate
  no neighbor 2001:DB8::2 activate
  no neighbor 2001:DB8::3 activate
  no neighbor 2001:DB8:14::4 activate
 exit-address-family
 !
 address-family ipv6
  neighbor 2001:DB8::2 activate
  neighbor 2001:DB8::2 next-hop-self
  neighbor 2001:DB8::3 activate
  neighbor 2001:DB8::2 next-hop-self
  neighbor 2001:DB8:14::4 activate
```

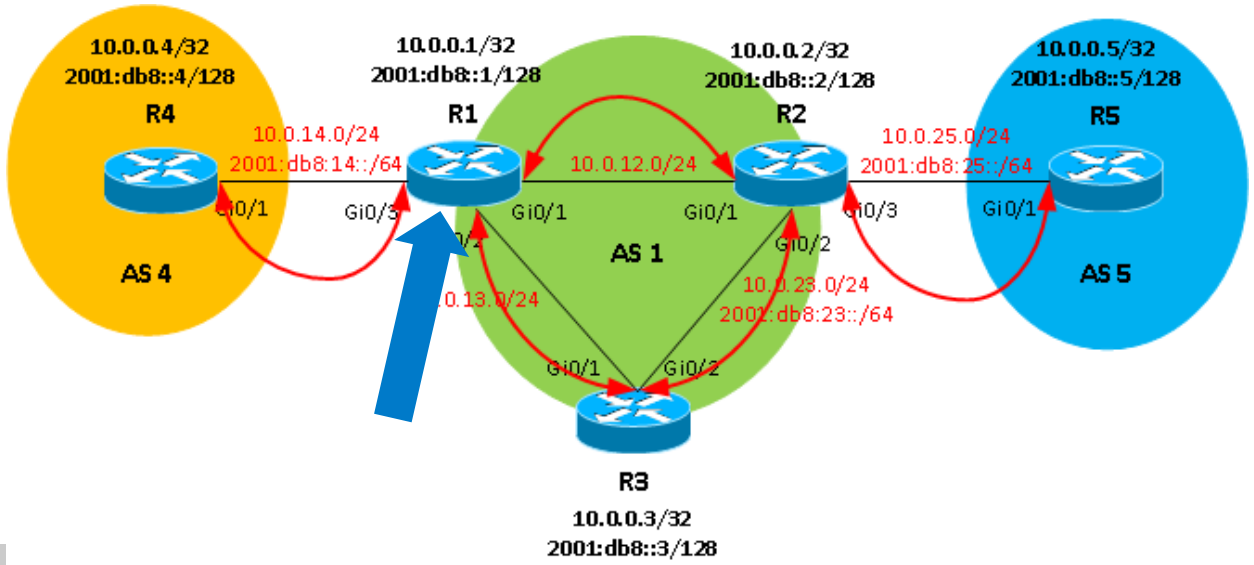


```
!
router bgp 4
 bgp router-id 10.0.0.4
 bgp log-neighbor-changes
 neighbor 10.0.14.1 remote-as 1
 neighbor 2001:DB8:14::1 remote-as 1
 !
 address-family ipv4
  network 10.0.0.4 mask 255.255.255.255
  neighbor 10.0.14.1 activate
  no neighbor 2001:DB8:14::1 activate
 exit-address-family
 !
 address-family ipv6
  network 2001:DB8::4/128
  neighbor 2001:DB8:14::1 activate
 exit-address-family
!
```

© 20

BGP Neighbours

```
R1#show bgp ipv6 unicast neighbors 2001:db8::2
BGP neighbor is 2001:DB8::2,  remote AS 1, internal link
  BGP version 4, remote router ID 10.0.0.2
  BGP state = Established, up for 00:16:28
....

 For address family: IPv6 Unicast
  Session: 2001:DB8::2
  BGP table version 3, neighbor version 3/0
  Output queue size : 0
  Index 1, Advertise bit 0
  1 update-group member
  NEXT_HOP is always this router for eBGP paths
  Slow-peer detection is disabled
  Slow-peer split-update-group dynamic is disabled
                                 Sent       Rcvd
  Prefix activity:               ----       ----
    Prefixes Current:               1          1 (Consumes 100 bytes)
    Prefixes Total:                 1          1
    Implicit Withdraw:              0          0
    Explicit Withdraw:              0          0
    Used as bestpath:             n/a          1
    Used as multipath:            n/a          0

…
Connection state is ESTAB, I/O status: 1, unread input bytes: 0
Connection is ECN Disabled, Mininum incoming TTL 0, Outgoing TTL 255
Local host: 2001:DB8::1, Local port: 179
Foreign host: 2001:DB8::2, Foreign port: 34782
Connection tableid (VRF): 0
....
R1#
```

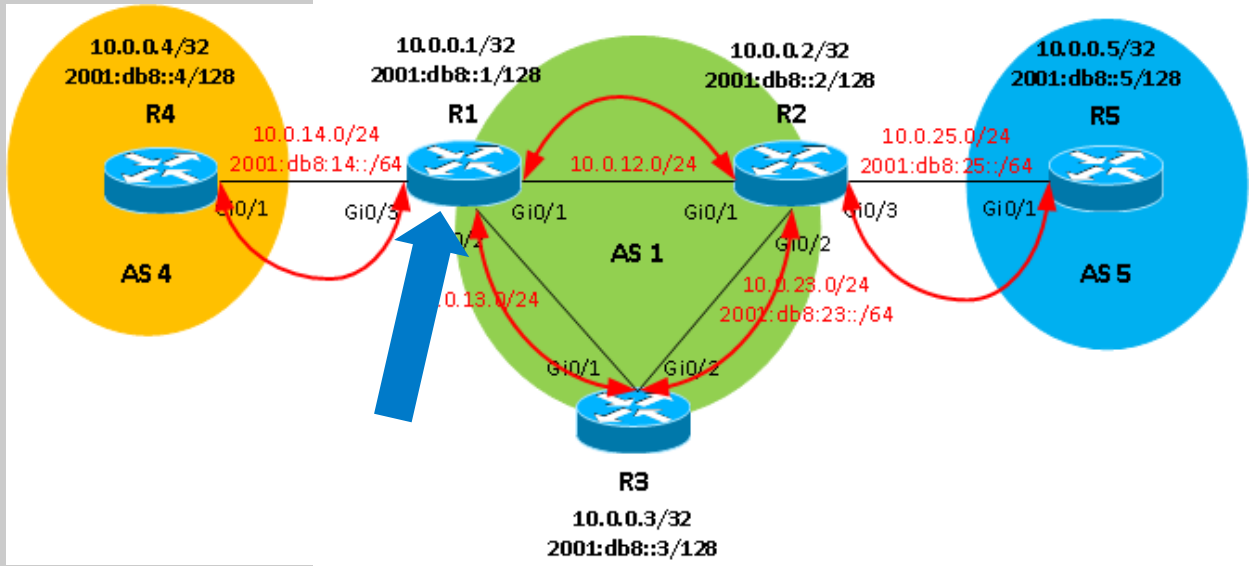

For Your Reference

© 2014 Cisco and/or its affiliates. All rights reserved.

BGP Prefix Table

```
R1#show bgp ipv6 unicast 2001:db8::5/128
BGP routing table entry for 2001:DB8::5/128, version 3
Paths: (1 available, best #1, table default)
  Advertised to update-groups:
     2
  Refresh Epoch 1
  5
    2001:DB8::2 (metric 1) from 2001:DB8::2 (10.0.0.2)
      Origin IGP, metric 0, localpref 100, valid, internal, best
      rx pathid: 0, tx pathid: 0x0
R1#
```

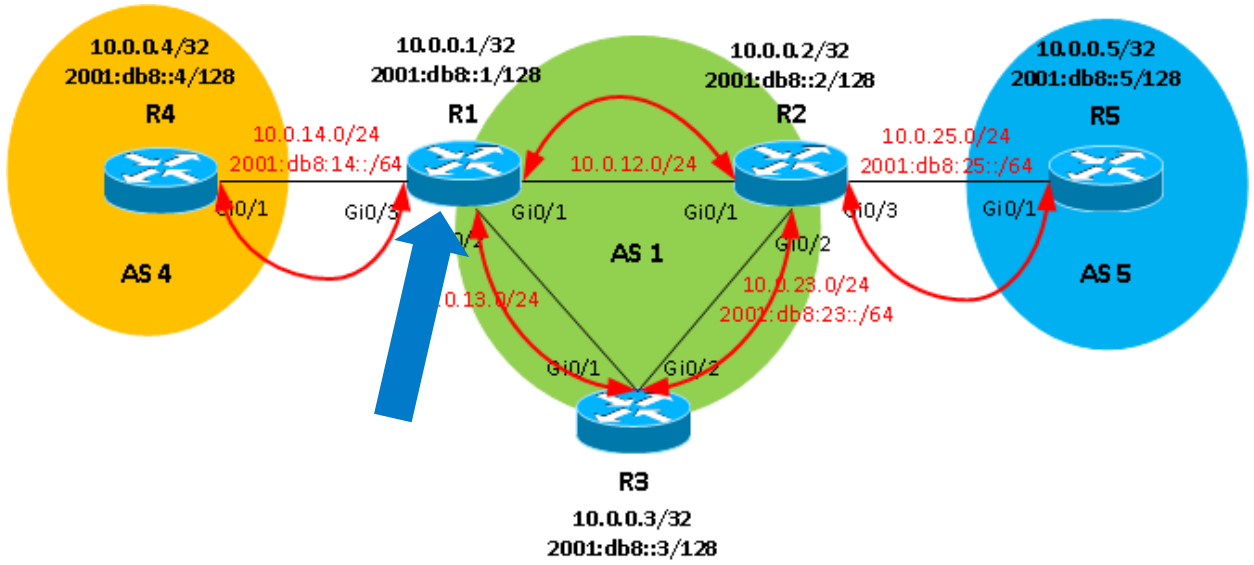


```
R1#show bgp ipv6 unicast 2001:db8::4/128
BGP routing table entry for 2001:DB8::4/128, version 2
Paths: (1 available, best #1, table default)
  Advertised to update-groups:
     1           3
  Refresh Epoch 1
  4
    2001:DB8:14::4 (FE80::F816:3EFF:FEE6:97BE) from 2001:DB8:14::4
(10.0.0.4)
      Origin IGP, metric 0, localpref 100, valid, external, best
      rx pathid: 0, tx pathid: 0x0
R1#
```

For Your Reference

Cisco live!

© 2014 Cisco and/or its affiliates. All rights reserved.

BGP IPv6 Link-Local Address Peering

```
R1#show run int gigabitEthernet 0/3
Building configuration...

Current configuration : 182 bytes
!
interface GigabitEthernet0/3
 ip address 10.0.14.1 255.255.255.0
 ....
 ipv6 address FE80:14::1 link-local
end

R1#
```

```
R4#show run int gigabitEthernet 0/1
Building configuration...

Current configuration : 182 bytes
!
interface GigabitEthernet0/1
 ip address 10.0.14.4 255.255.255.0
 ....
ipv6 address FE80:14::4 link-local
end

R4#
```

```
ck0
remote-as 1
0/1 activate
  address-family ipv6
   network 2001:DB8::4/128
   neighbor FE80:14::1%GigabitEthernet0/1 activate
  exit-address-family
 !
 exit-address-family
!
```

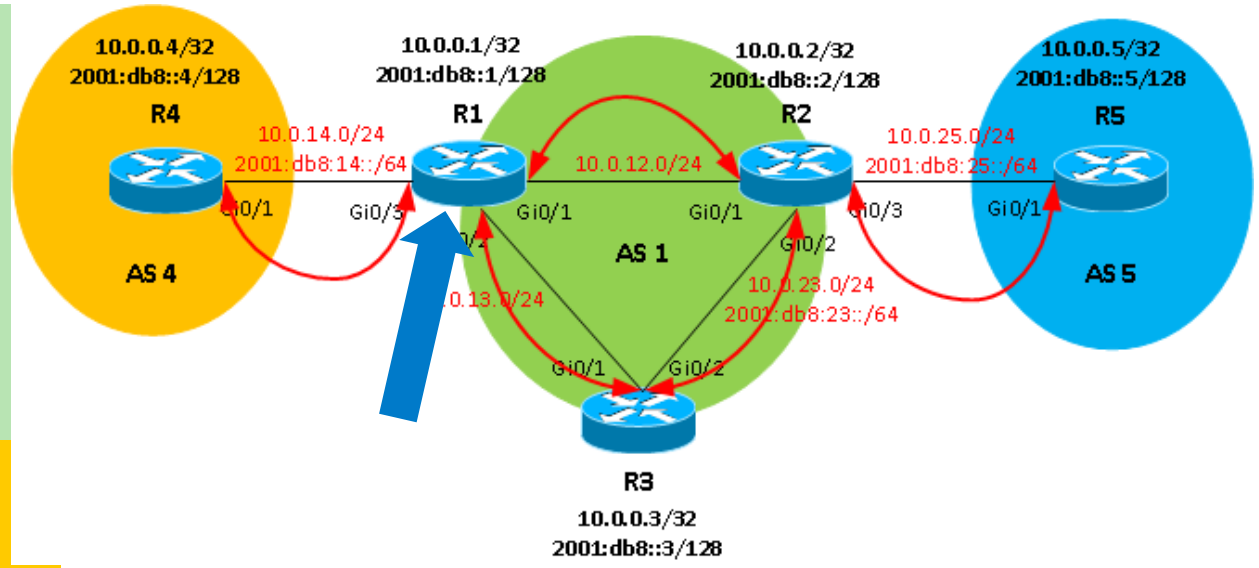


```
R1#show bgp ipv6 unicast 2001:db8::4/128
BGP routing table entry for 2001:DB8::4/128, version 11
Paths: (1 available, best #1, table default)
  Advertised to update-groups:
     1           3
  Refresh Epoch 1
  4
    FE80:14::4 (FE80:14::4) from FE80:14::4%GigabitEthernet0/3 (10.0.0.4)
      Origin IGP, metric 0, localpref 100, valid, external, best
      rx pathid: 0, tx pathid: 0x0
R1#
```

EIGRP for IPv6

EIGRP for IPv6

Implementation	Three new TLVs: •IPv6_REQUEST_TYPE (0X0401) •IPv6_METRIC_TYPE (0X0402) •IPv6_EXTERIOR_TYPE (0X0403) New Protocol Dependent Module (PDM) 32 bit Router ID which must be explicitly configured if no IPv4 address Same Protocol Number, Metrics
Important Differences	Hellos are sourced from the link-local address and destined to FF02::A (all EIGRP routers); this means that neighbors do not have to share the same global prefix (with the exception of explicitly specified neighbors where traffic is unicasted) Automatic summarization disabled by default and not configurable for IPv6 (same as in IPv4 now – CSCso20666)
Notes	The IPv6 EIGRP process must be enabled explicitly ("no shutdown" under the process) in older IOS version. In current versions, the default is that IPv6 EIGRP is enabled (CSCso90068) EIGRP now going through the IETF process as an informational RFC: http://tools.ietf.org/html/draft-savage-eigrp-01

© 2014 Cisco and/or its affiliates. All rights reserved.

EIGRP for IPv6 and IPv6 Summary

	EIGRP – IPv4	EIGRP – IPv6
Type	Distance vector	Distance vector
Metric	Composite (in practice: bandwidth and delay)	Composite (in practice: bandwidth and delay)
Loop prevention	DUAL, split horizon	DUAL, split horizon
Administrative distance	5 (sum.), 90 (int.), 170 (ext.)	5 (sum.), 90 (int.), 170 (ext.)
L3 transport S/D	IPv4 unicast / 224.0.0.10	IPv6 LL unicast / FF02::A
L4 transport	IP 88, RTP (reliable multicast)	Next Header 88, RTP

```
interface Loopback0
 ip address 10.10.0.1 255.255.255.255
!
interface Ethernet0/0
 ip address 10.0.0.1 255.255.255.0
!
router eigrp 1
 network 10.0.0.0
 passive-interface Loopback0
```

```
ipv6 unicast-routing
!
interface Loopback0
 ipv6 address 2001:DB8:1000::1/128
 ipv6 eigrp 1
!
interface Ethernet0/0
 ipv6 address 2001:DB8::1/64
 ipv6 eigrp 1
!
ipv6 router eigrp 1
 passive-interface Loopback0
 eigrp router-id 10.10.10.1
```

For Your Reference

© 2014 Cisco and/or its affiliates. All rights reserved.

EIGRP Sample Topology

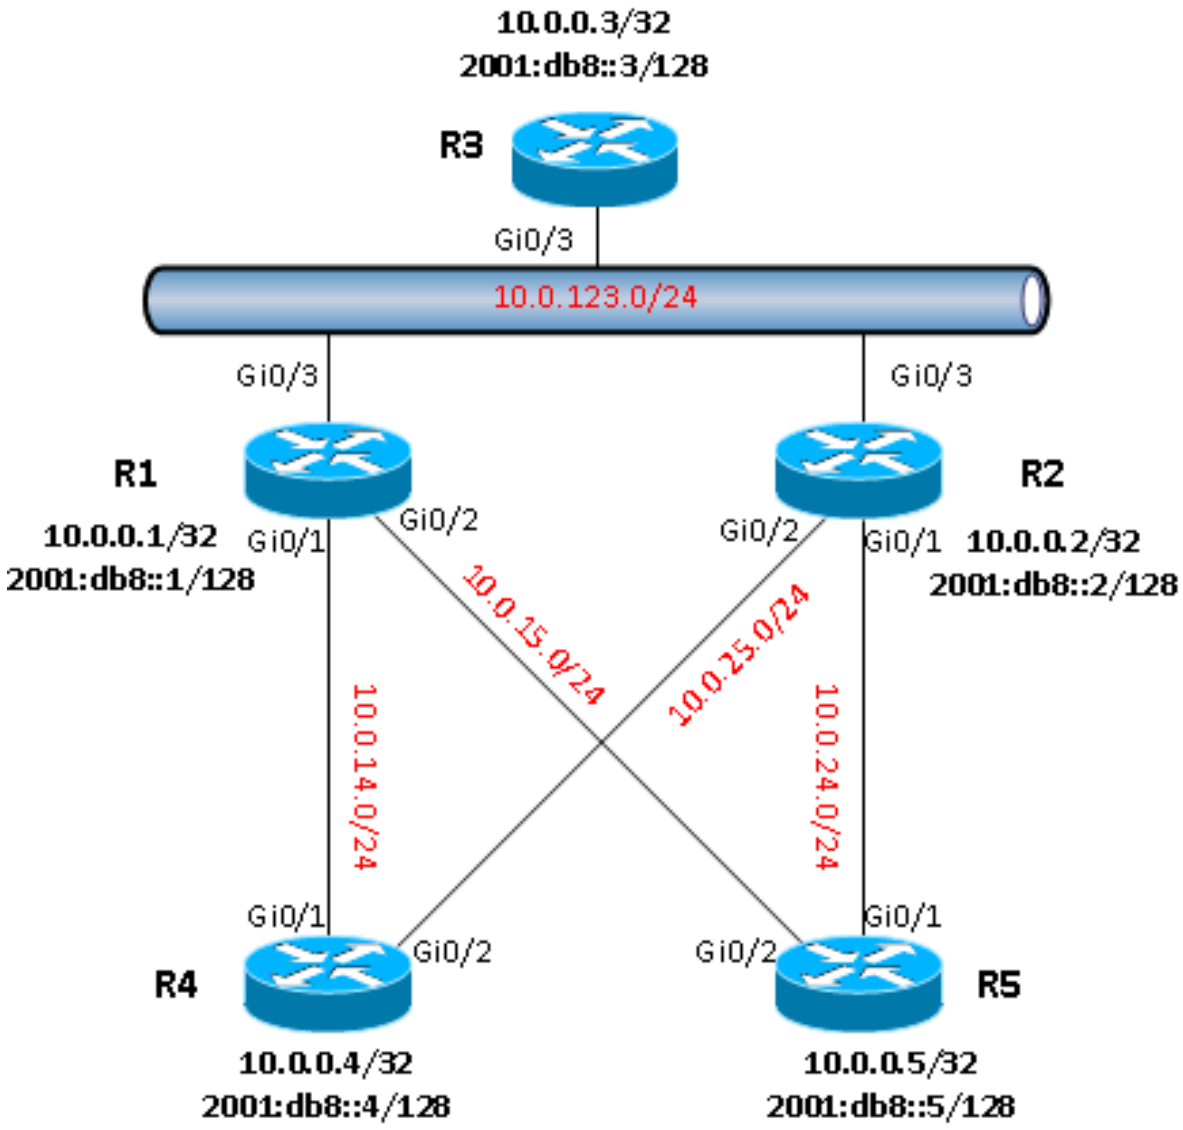

© 2014 Cisco and/or its affiliates. All rights reserved.

EIGRP Configuration

```
!
ipv6 unicast-routing
!
interface Loopback0
 ipv6 address 2001:DB8::1/128
 ipv6 eigrp 1
!
interface GigabitEthernet0/1
 delay 10
 ipv6 enable
 ipv6 eigrp 1
!
interface GigabitEthernet0/2
 delay 5
 ipv6 enable
 ipv6 eigrp 1
!
interface GigabitEthernet0/3
 delay 1
 ipv6 enable
 ipv6 eigrp 1
!
ipv6 router eigrp 1
 passive-interface Loopback0
 eigrp router-id 10.0.0.1
!
```

```
!
interface Loopback0
 ip address 10.0.0.1 255.255.255.255
!
interface GigabitEthernet0/1
 ip address 10.0.14.1 255.255.255.0
 delay 10
!
interface GigabitEthernet0/2
 ip address 10.0.15.1 255.255.255.0
 delay 5
!
interface GigabitEthernet0/3
 ip address 10.0.123.1 255.255.255.0
 delay 1
!
router eigrp 1
 network 10.0.0.1 0.0.0.0
 network 10.0.14.0 0.0.0.255
 network 10.0.15.0 0.0.0.255
 network 10.0.123.0 0.0.0.255
 passive-interface Loopback0
 eigrp router-id 10.0.0.1
!
```

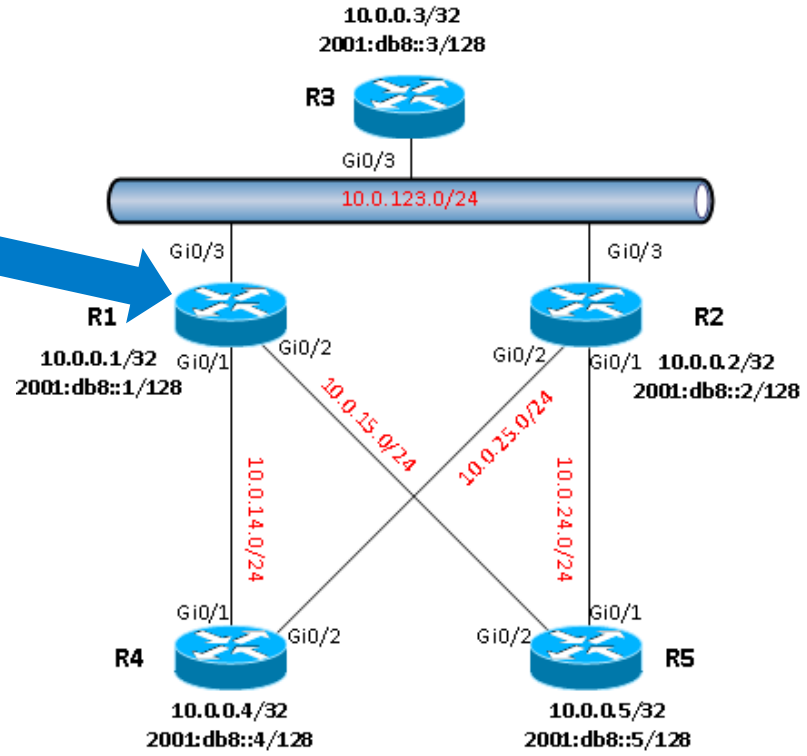

© 2014 Cisco and/or its affiliates. All rights reserved.

EIGRP Interfaces

```
R1#sh ipv6 interface gigabitEthernet 0/1
GigabitEthernet0/1 is up, line protocol is up
  IPv6 is enabled, link-local address is FE80::F816:3EFF:FEC6:DE10
  No Virtual link-local address(es):
  No global unicast address is configured
  Joined group address(es):
```

```
R1#show eigrp address-family ipv6 interfaces detail
EIGRP-IPv6 Interfaces for AS(1)
                              Xmit Queue   PeerQ        Mean   Pacing Time   Multicast
Pending
Interface              Peers  Un/Reliable  Un/Reliable  SRTT   Un/Reliable   Flow Timer
Routes
Gi0/1                    1        0/0         0/0       162       0/0           812          0
  Hello-interval is 5, Hold-time is 15
  Split-horizon is enabled
  Next xmit serial <none>
  Packetized sent/expedited: 25/1
  Hello's sent/expedited: 43344/4
  Un/reliable mcasts: 0/21  Un/reliable ucasts: 36/11
  Mcast exceptions: 0  CR packets: 0  ACKs suppressed: 0
  Retransmissions sent: 6  Out-of-sequence rcvd: 4
  Interface has all stub peers
  Topology-ids on interface - 0
  Authentication mode is not set
....
```

```
  Retransmissions sent: 5  Out-of-sequence rcvd: 2
  Interface has all stub peers
  Topology-ids on interface - 0
  Authentication mode is not set
....
```

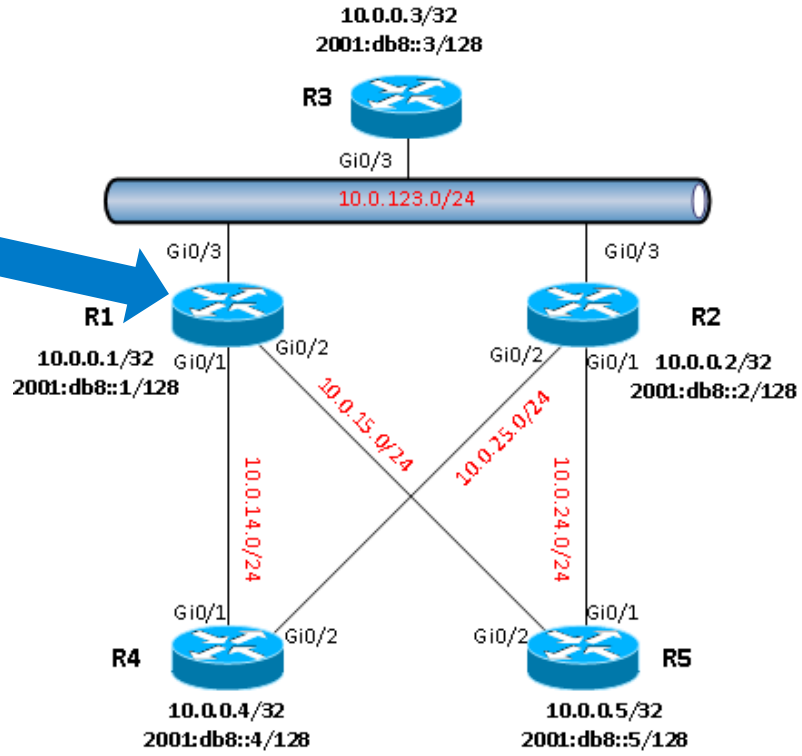

© 2014 Cisco and/or its affiliates. All rights reserved.

EIGRP Neigbours

```
R1#show eigrp address-family ipv6 neighbors detail
EIGRP-IPv6 Neighbors for AS(1)
H   Address                 Interface              Hold Uptime   SRTT   RTO  Q   Seq
                                                   (sec)         (ms)       Cnt Num
3   Link-local address:     Gi0/2                   10 03:41:14   84   504  0   74
    FE80::F816:3EFF:FE73:39F
   Version 12.0/2.0, Retrans: 0, Retries: 0, Prefixes: 1
   Topology-ids from peer - 0
   Stub
   Supp
2   Linl
    FE8(
   Vers
   Topo
   Stub
   Supp
1   Linl
    FE8(
   Vers
   Topo
0   Linl
    FE8(
   Vers
   Topo
R1#
```

```
R1#show eigrp address-family ipv4 neighbors detail
EIGRP-IPv4 Neighbors for AS(1)
H   Address                 Interface              Hold Uptime   SRTT   RTO  Q   Seq
                                                   (sec)         (ms)       Cnt Num
3   10.0.15.5               Gi0/2                   10 03:09:09 1598   5000  0   19
   Version 12.0/2.0, Retrans: 1, Retries: 0, Prefixes: 2
   Topology-ids from peer - 0
   Stub Peer Advertising (CONNECTED ) Routes
   Suppressing queries
0   10.0.14.4               Gi0/1                   12 03:09:23  817   4902  0   24
   Version 12.0/2.0, Retrans: 1, Retries: 0, Prefixes: 2
   Topology-ids from peer - 0
   Stub Peer Advertising (CONNECTED ) Routes
   Suppressing queries
2   10.0.123.2              Gi0/3                   13 03:16:20    9   100  0   32
   Version 12.0/2.0, Retrans: 3, Retries: 0, Prefixes: 4
   Topology-ids from peer - 0
1   10.0.123.3              Gi0/3                   11 03:17:43    9   100  0   13
   Version 12.0/2.0, Retrans: 4, Retries: 0, Prefixes: 1
   Topology-ids from peer - 0
R1#
```

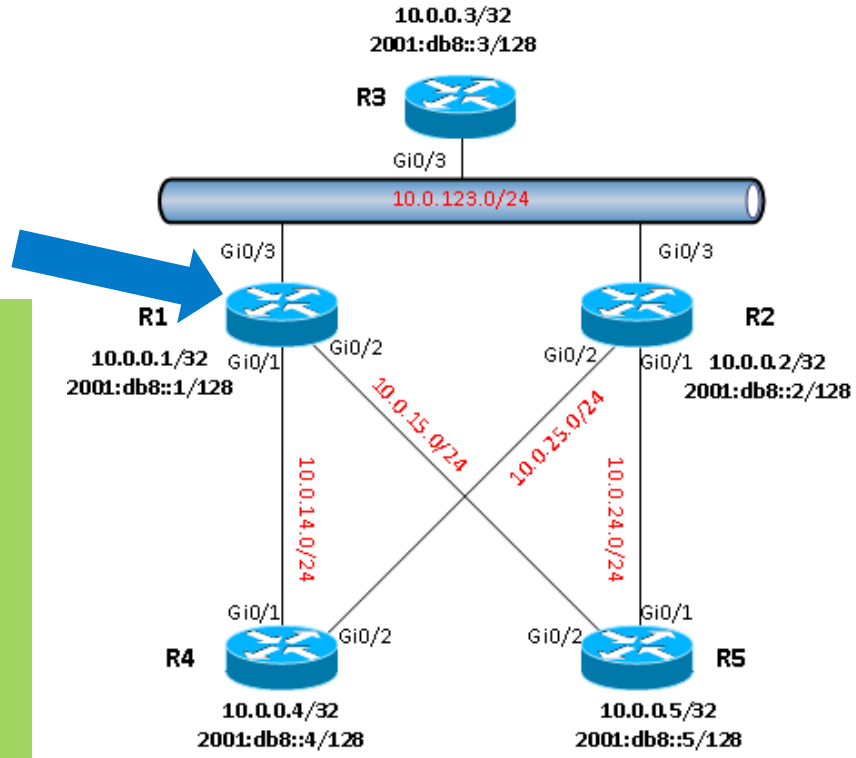

© 2014 Cisco and/or its affiliates. All rights reserved.

EIGRP Topology Table

```
R4#show eigrp address-family ipv6 topology
EIGRP-IPv6 Topology Table for AS(1)/ID(10.0.0.4)
Codes: P - Passive, A - Active, U - Update, Q - Query, R - Reply,
       r - reply Status, s - sia Status

P 2001:DB8::
        via
GigabitEther
        via
GigabitEther
P 2001:DB8::
        via
GigabitEther
        via
GigabitEther
P 2001:DB8::
        via
GigabitEther
        via
GigabitEther
P 2001:DB8::
        via (
P 2001:DB8::
        via
GigabitEther
        via
GigabitEther
```

```
R4#
```

```
R4#show eigrp address-family ipv4 topology
EIGRP-IPv4 Topology Table for AS(1)/ID(10.0.0.4)
Codes: P - Passive, A - Active, U - Update, Q - Query, R - Reply,
       r - reply Status, s - sia Status

P 10.0.25.0/24, 1 successors, FD is 6400
        via 10.0.24.2 (6400/5120), GigabitEthernet0/2
        via 10.0.14.1 (7936/5376), GigabitEthernet0/1
P 10.0.15.0/24, 1 successors, FD is 5376
        via 10.0.24.2 (5376/4096), GigabitEthernet0/2
        via 10.0.14.1 (6400/3840), GigabitEthernet0/1
P 10.0.0.3/32, 1 successors, FD is 132096
        via 10.0.24.2 (132096/130816), GigabitEthernet0/2
        via 10.0.14.1 (133376/130816), GigabitEthernet0/1
P 10.0.0.2/32, 1 successors, FD is 131840
        via 10.0.24.2 (131840/128256), GigabitEthernet0/2
        via 10.0.14.1 (133376/130816), GigabitEthernet0/1
P 10.0.14.0/24, 1 successors, FD is 5120
        via Connected, GigabitEthernet0/1
P 10.0.0.4/32, 1 successors, FD is 128256
        via Connected, Loopback0
P 10.0.0.5/32, 1 successors, FD is 133376
        via 10.0.24.2 (133376/132096), GigabitEthernet0/2
        via 10.0.14.1 (134400/131840), GigabitEthernet0/1
P 10.0.0.1/32, 1 successors, FD is 132096
        via 10.0.24.2 (132096/130816), GigabitEthernet0/2
        via 10.0.14.1 (133120/128256), GigabitEthernet0/1
. . .
```

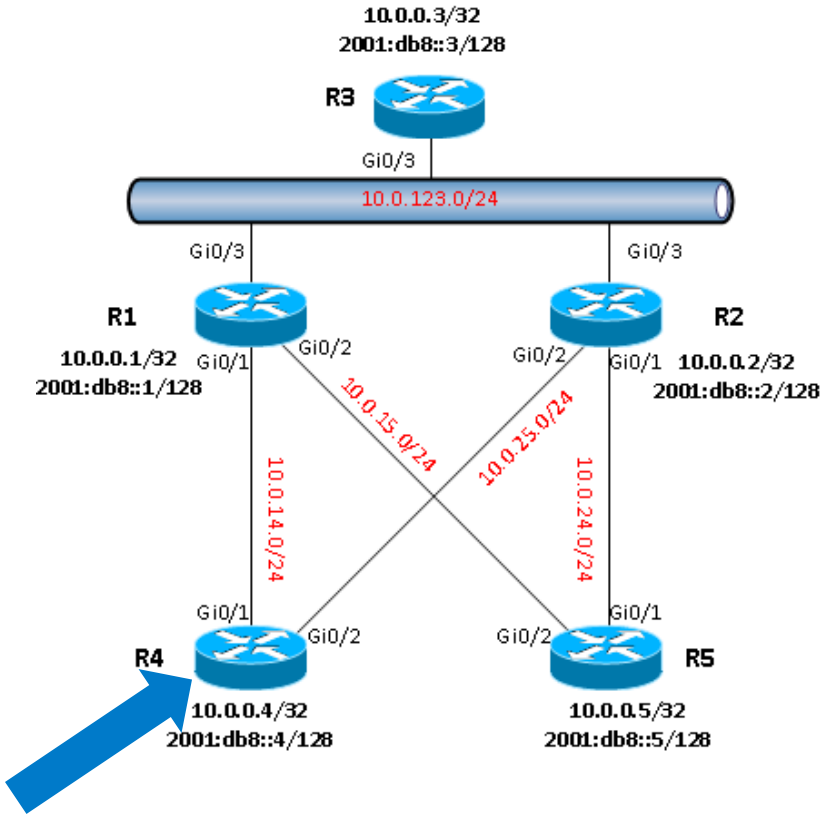

EIGRP RIB

```
R4#show ipv6 route eigrp
IPv6 Routing Table - default - 6 entries
Codes: C - Connected, L - Local, S - Static, U - Per-user Static route
       B - BGP, HA - Home Agent, MR - Mobile Router, R - RIP
       H - NHRP, I1 - ISIS L1, I2 - ISIS L2, IA - ISIS interarea
       IS - ISIS summary, D - EIGRP, EX - EIGRP external, NM - NEMO
       ND - ND Default, NDp - ND Prefix, DCE - Destination, NDr -
Redirect
       O - OSPF Intra, OI - OSPF Inter, OE1 - OSPF ext 1, OE2 - OSPF ext
2
       ON1 - OSP
       ld - LISP
D   2001:DB8::1/
     via FE80::
D   2001:DB8::2/
     via FE80::
D   2001:DB8::3/
     via FE80::
D   2001:DB8::5/
     via FE80::
R4#
```

```
R4#show ip route eigrp
Codes: L - local, C - connected, S - static, R - RIP, M - mobile, B - BGP
       D - EIGRP, EX - EIGRP external, O - OSPF, IA - OSPF inter area
       N1 - OSPF NSSA external type 1, N2 - OSPF NSSA external type 2
       E1 - OSPF external type 1, E2 - OSPF external type 2
       i - IS-IS, su - IS-IS summary, L1 - IS-IS level-1, L2 - IS-IS level-2
       ia - IS-IS inter area, * - candidate default, U - per-user static
route
       o - ODR, P - periodic downloaded static route, H - NHRP, l - LISP
       a - application route
       + - replicated route, % - next hop override

Gateway of last resort is 10.11.12.1 to network 0.0.0.0

      10.0.0.0/8 is variably subnetted, 14 subnets, 2 masks
D        10.0.0.1/32 [90/132096] via 10.0.24.2, 03:18:51, GigabitEthernet0/2
D        10.0.0.2/32 [90/131840] via 10.0.24.2, 03:18:51, GigabitEthernet0/2
D        10.0.0.3/32 [90/132096] via 10.0.24.2, 03:18:51, GigabitEthernet0/2
D        10.0.0.5/32 [90/133376] via 10.0.24.2, 03:18:37, GigabitEthernet0/2
D        10.0.15.0/24 [90/5376] via 10.0.24.2, 03:18:51, GigabitEthernet0/2
D        10.0.25.0/24 [90/6400] via 10.0.24.2, 03:18:51, GigabitEthernet0/2
D        10.0.123.0/24 [90/4096] via 10.0.24.2, 03:18:51, GigabitEthernet0/2
R4#
```

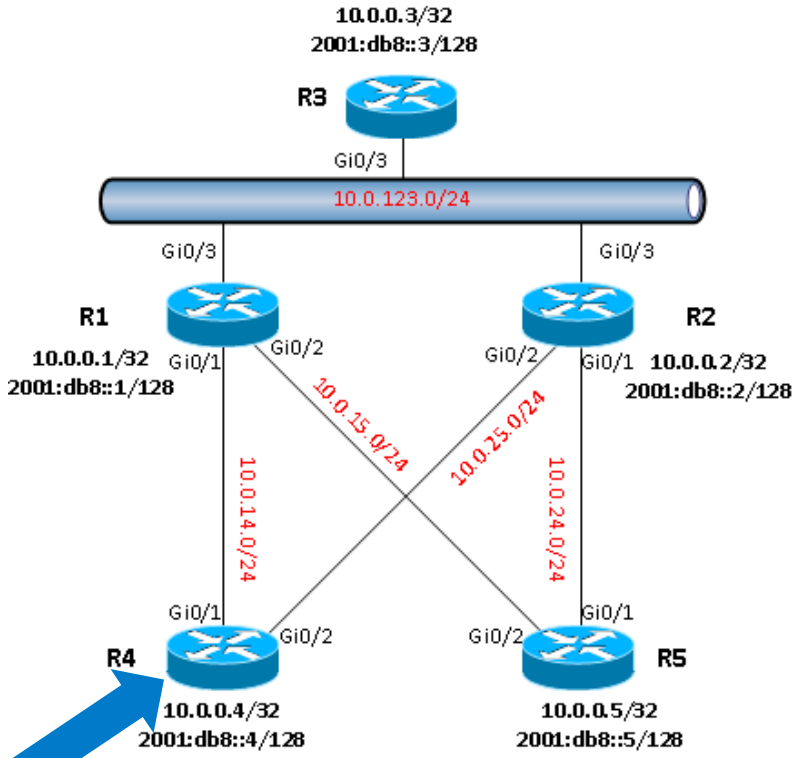

© 2014 Cisco and/or its affiliates. All rights reserved.
Cisco Public

Co-existence of IPv4 and IPv6 Routing Protocols

IPv4 and IPv6 Co-existence

	Single Process / Single Topology	Single Process / Multi Topology	Multi Process / Multi Topology
Protocols	IS-IS ST	IS-IS MT	OSPFv2 + OSPFv3 EIGRP + EIGRPv6
IP topologies	Single (IPv4+IPv6) Congruent	Multiple Non-congruent	Multiple Non-congruent
Flooding + router/network resources	Common	Common	Multiple protocol instances on given link
SPF	Single	Multiple	Multiple (OSPF)
LS databases / topology tables	Single Large	Single Large	Multiple
Control plane For Your Reference	- Common - Less resource intensive - More deterministic IPv4/ IPv6 co-existence	- More separation - Protocol-specific optimization possible - More resource intensive	- Clear separation - Protocol-specific optimization possible - More resource intensive

© 2014 Cisco and/or its affiliates. All rights reserved.

The IPv4 / IPv6 Co-existence Twist

- IPv6 IGP impact on the IPv4 IGP convergence
- Aggressive timers on both IGPs will highlight competition for resources
- Is parity necessary from day 1?

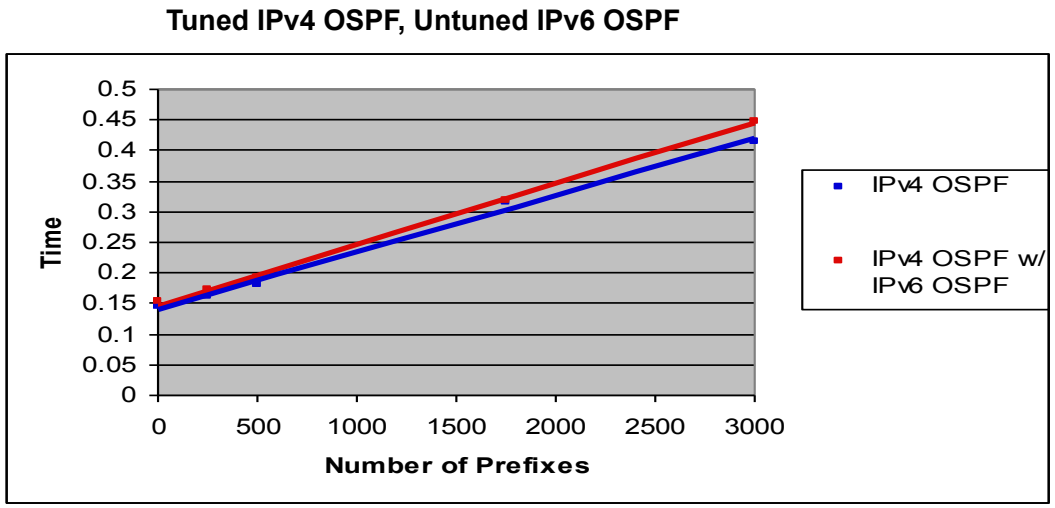

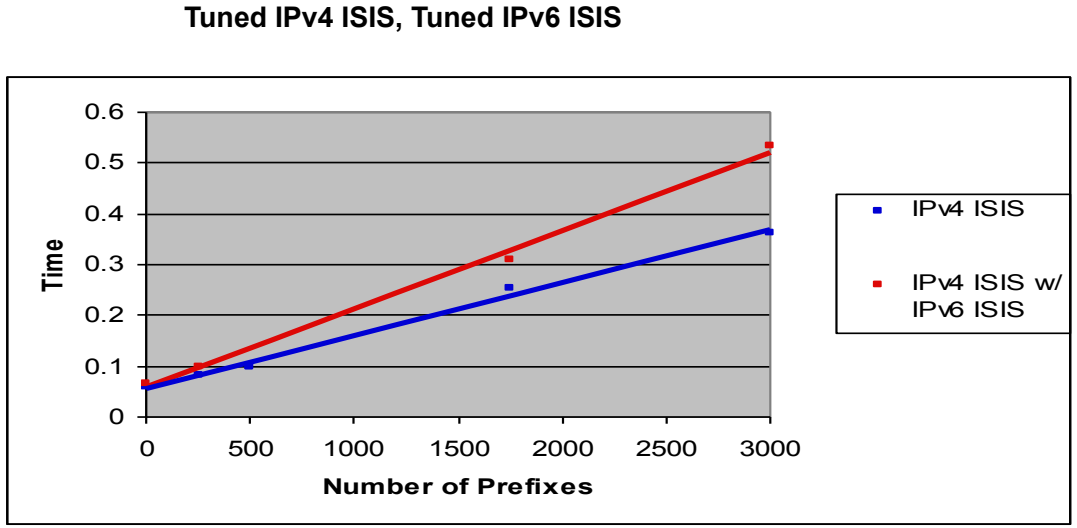

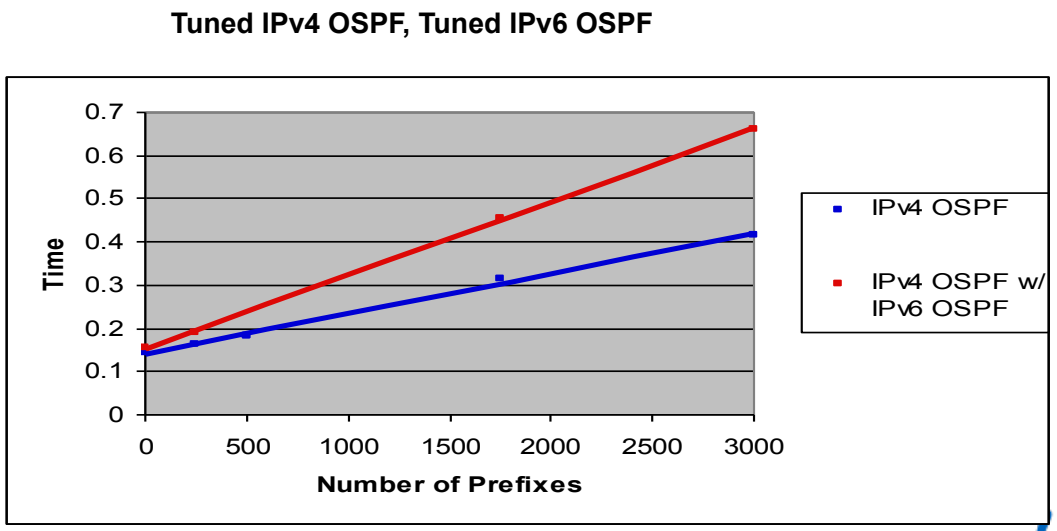

© 2014 Cisco and/or its affiliates. All rights reserved.

Other Considerations

Other Considerations

- Routing protocol features !
  - Check on CCO
- Which routing protocol is your team comfortable with?
  - Design and engineering
  - Operations
  - Support
  - I would suspect 99% of customer to stay within the same “routing protocol family”

© 2014 Cisco and/or its affiliates. All rights reserved.

IPv6 Fast Convergence and GR Summary

	OSPFv3	IS-IS for IPv6	EIGRP for IPv6
LSA/LSP generation exponential backoff throttling	Y	Y	N/A
SPF execution exponential backoff throttling	Y	Y	N/A
Bidirectional Forwarding Detection (BFD)	Y	Y (through IPv4 BFD)	Y (CSCtk04807)
Graceful Restart (GR)	Y	Y	Y

For Your Reference

Cisco live!

© 2014 Cisco and/or its affiliates. All rights reserved.

Conclusions

IPv4 / IPv6 IGP Implementation Summary (IOS)

Implementation	Multi Topology	Multi Instance	Multi Process
Processes	Single	Single	Multiple
Topologies / SPFs	Multiple	Multiple	Multiple
Adjacencies	Single	Single	Multiple
EIGRPv6 / EIGRP			Yes (same protocols, different PDMs)
IS-IS	Yes. (ST IS-IS requires identical topology / SPF for IPv4 and IPv6)		
OSPFv3 / OSPFv2			Yes (different protocols)
OSPFv3 with AF support (RFC 5838)		Yes (IPv4 and IPv6 support)	

Conclusions

- A dual-stack architecture requires an IPv6 routing protocol to be deployed alongside the existing IPv4 routing protocol
- In many aspects, the IPv6 routing protocols are very similar to their IPv4 counterparts
  - You don't have to start from scratch!
  - Some IPv6 routing protocol specificities (e.g. the use of Link-Local NH addresses, multi-topology implementations) need to be understood
  - OSPFv3 has been significantly re-designed compared to OSPFv2
- The criteria to choose one IPv6 routing protocol vs. another one (e.g. link state vs. DUAL) are identical to the ones for choosing an IPv4 routing protocol
- Design and deployment considerations for IPv6 routing protocols are similar to their IPv4 counterparts
- Co-existence of IPv4 and IPv6 routing protocols needs to be considered
  - Impact on hardware resources – CPU and memory

© 2014 Cisco and/or its affiliates. All rights reserved.

CISCO™